

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 987 861 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
22.03.2000 Bulletin 2000/12

(51) Int Cl.7: H04L 12/56, H04Q 11/04

(21) Application number: 99307100.0

(22) Date of filing: 07.09.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• **Chrin, Christopher James**
Lisle, Illinois 60532 (US)
• **Zola, Meyer Joseph**
Lisle, Illinois 60532 (US)

(30) Priority: 18.09.1998 US 101030 P
25.03.1999 US 276187


(74) Representative:
Watts, Christopher Malcolm Kelway, Dr. et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(71) Applicant: LUCENT TECHNOLOGIES INC.
Murray Hill, New Jersey 07974-0636 (US)

(54) Flexible telecommunications switching network

(57) Apparatus, and a method for flexibly switching telecommunication signals. A plurality of input bit streams are connected to a telecommunications switching network fabric element comprising a microprocessor system including memory. The microprocessor system, under program control, performs switching and protocol conversion functions on the input streams in order to generate output streams. Advantageously, a single ele-

ment is able to concurrently switch input signals in a variety of protocols, including circuit switching protocols, such as Pulse code Modulation (PCM), and packet switching protocols, such as Asynchronous Transfer Mode (ATM) protocols and Internet Protocol, (IP). Where desirable, the microprocessor system can also control the protocol conversion of input signals in one protocol to output signals in another protocol.



EP 0 987 861 A2

Description

Technical Field:

[0001] This invention relates to digital telecommunications switching network fabrics.

Problem:

[0002] Telecommunications switching network fabrics are the means by which individual telecommunications messages and/or circuit connections are routed from an input of a switching system to an output of such a switching system. The input and/or output can be connected either to another switching system, or to a source/destination of the messages or connections.

[0003] Largely as a result of the ubiquitous use of optical fiber transmission systems to interconnect switching systems, the switching of modern switching network fabrics is almost entirely performed in the digital mode. For standard telephone conversations, this switching is largely performed by switching pulse code modulation (PCM) signals through time slot interchange (TSI), and time multiplexed (TMS) switches. For the increasing amount of data traffic, packet switches are required. Such switches take a packet from an input data stream, examine and in general, alter, a header identifying the path that the packet is to take, and route the packet to the output indicated by the packet header. Two of the most common types of data inputs are those in the asynchronous transfer mode (ATM), and those in the Internet Protocol (IP).

[0004] A problem of the prior art is that the switching network fabrics are in general designed to handle only one type of input and output stream, and that special conversion equipment where necessary, and multiple switching network fabrics are required when multiple types of input and output streams are encountered. One exception to the above statement is that disclosed in U. S. Patent 5,345,446, which discloses arrangements for converting input/output PCM streams into ATM format, and switching the converted as well as the native ATM inputs to a switching system through a common ATM fabric. This approach requires expensive conversion equipment, and is less efficient if the facilities interconnecting switching systems are largely, or predominantly PCM facilities.

Solution:

[0005] The above problem is solved and an advance is made over the prior art in accordance with invention, wherein the input and output streams of a switching network fabric are performed under the control of, and essentially within, a microprocessor. The program of the microprocessor determines how different types of input streams, such as PCM, ATM, or IP are to be switched, and accesses all necessary internal or external memory

to select an output stream, and for the case of packet streams, how to alter the header of each received packet. Advantageously, any mix of different types of input stream protocols can be switched, and where necessary, converted, under the control of a single microprocessor entity. If the traffic on one particular input stream changes from one type of protocol to another, the input stream need not be reconnected to a different switching entity; instead the microprocessor is notified of the new type of protocol for this input stream, and performs its switching accordingly.

[0006] In accordance with one aspect of this invention, the microprocessor switching entity performs conversions, e.g., between PCM and ATM, as necessary to meet the requirements of an output transmission facility. Advantageously, no separate equipment and separate routing is required.

[0007] In accordance with one preferred embodiment, a RISC communications processor such as the PowerPC® manufactured by the Motorola Corporation, is used as the microprocessor of the switching network fabric. A three hundred MHZ processor, such as the EC 603e of this type can, for example, handle up to 192 PCM streams each consisting of 32 time slots at a bit rate of 2.048 Mbits per second. The capacity is likely to be less if the microprocessor handles packet as well as PCM traffic. On the other hand, if a PowerPC or equivalent processor, is manufactured with a higher clock rate, such a microprocessor can handle more traffic. In one Application, the Input/Output bandwidth and/or the cache size of the processor may be limiting. Advantageously, today's technology will support a very substantial size switching network fabric unit (element).

[0008] In accordance with one arrangement for creating a larger network, a plurality of microprocessor fabric units have their inputs connected in parallel, but each is connected to a different set of output streams. The output streams in some cases, to transmission facilities, and in other cases, to another microprocessor network fabric unit.

[0009] In one embodiment, the microprocessor elements generate a bit stream in an output protocol that is different from the input protocol; the output bit stream can then be further switched by other units.

[0010] In accordance with another aspect of Applicants' invention, groups of circuits, or packets, can be switched using this arrangement. Advantageously, such an arrangement can be used efficiently to implement a facilities switch of a type which switches groups of circuits and which responds not to individual call set-up requests, but to group set-up requests, usually from some operating support system.

Brief Description of The Drawing(S):

[0011]

Figure 1 is a block diagram of a microprocessor

based time slot interchange, (TSI) module;

Figure 2 is a block diagram of the internal microprocessor architecture of the TSI module;

Figure 3 is a block diagram of memory and buffer layouts for the TSI module;

Figure 4 is a flow diagram of a program for controlling the TSI module;

Figure 5 shows modifications of the program of Figure 4, for the case in which the multiple sub-frames are encountered between each frame synchronizing pulse.

Figure 6 illustrates the program for controlling the microprocessor, acting as a time multiplexed switch;

Figure 7 shows the program for a TSI, wherein groups of time slots are bundled and switched as a bundle;

Figure 8 illustrates a cell header for an asynchronous transport method (ATM) cell;

Figure 9 illustrates the basic operation of the processing of an ATM cell;

Figure 10 is a programmer's data model, illustrating the layout of memory for a microprocessor acting as an ATM switch;

Figure 11 is a flow diagram, illustrating the operation of the input processing of a microprocessor acting as an ATM switch;

Figure 12 illustrates the processing of the output queue of the microprocessor program for an ATM switch;

Figure 13 illustrates the process of processing ATM cells to output links of a microprocessor acting as an ATM switch; and

Figure 14 is a block diagram, illustrating the arrangement for expanding the size of a microprocessor control switch through replication of microprocessor complexes.

Detailed Description:

[0012] This specification describes an arrangement for, and method of implementing multiple hardware functionality by appropriate software on a Reduced Instruction Set Computer, (RISC) microprocessor. Although this is described in terms of a RISC microproc-

essor, other types of microprocessor implementations, (e.g., Complex Instruction Set Computer), (CISC) can also be used. Multiple functions can reside on the same microprocessor simultaneously, or only a single function can be provided. The determination of the type of treatment each input and/or output shall receive, (e.g., Time Slot Interchange (TSI), Time Multiplexed Switch (TMS), Cross-connect (XCON), Asynchronous Transfer Mode (ATM) Switch, Internet Protocol (IP) Router, Dynamic Synchronous Transfer Mode (DTM), Frame Relay (FR) switch, etc.), is determined by, and can be reconfigured by, software control. Conversion between input to output formats can also be provided, e.g., a circuit switched PCM format can be converted to/from ATM format. A method is identified for using multiple microprocessors for building large configurations for applications which do not fit on a single microprocessor.

[0013] Other advantages include:

1. Little or no VLSI Required - faster time to market (eliminates development).
2. Microprocessor Self Test - (reduces investment in chip and board test tools).
3. Follows Moore's Law Technology Curve Directly - (graceful evolution).
4. Core Architecture Usable by Multiple Applications.
5. Results in Lower Development Effort (contributed to by all of the above).

[0014] Among the basic techniques used are the following:

1. Input bit streams are clocked into serial in, parallel out shift registers and then read out in parallel onto the microprocessor data bus under control of the microprocessor address bus.
2. This data, is then stored in the internal microprocessor cache memory which may be expanded by a level 2 cache on or outside the microprocessor chip, and/or an auxiliary memory outside the microprocessor chip, and is manipulated under stored program control in such a way as to provide the desired switching function.
3. The more frequently used sections of the stored program are, advantageously, stored in caches.
4. The resulting outputs are read out in parallel to parallel in, serial out shift registers, and then clocked out onto serial bit streams.

[0015] Figure 1 is a block diagram of the basic system

architecture. The heart of the system is a Microprocessor containing on board program and data caches and external Input/Output consisting of serial to parallel shift registers as Input Buffers, and parallel to serial shift registers as Output Buffers. An Input/Output decoder under the control of the Microprocessor selects the Input Buffer which puts its data on the jData Bus or the Output Buffer, which reads the data from the data bus. The Control Register is used to receive and transmit control messages from/to the outside world. The external Memory is used to store back-up and maintenance code, as well as data structures which are too large to fit into the cache. Microprocessor software is used to provide switching functionality which has traditionally been provided by hardware implementation. Thus, multiple diverse functions can be provided concurrently by a single microprocessor architecture. By changing the resident software, different sets of functions can be provided.

[0016] Figure 1 shows the input and output data streams of a time slot interchange unit 100 in accordance with Applicants' invention. In one preferred embodiment of Applicants' invention, a 300 MHZ PowerPC, EC 603e, can switch 192 serial input and output streams, each consisting of 32 time slots at a bit rate of 2.048 Mbits per second. The input comprises n serial input streams, stream zero being connected to input buffer 101, . . . , and serial input stream, n - 1, being connected to input buffer 102. The first input stream is collected in a shift register of input buffer 101, and then transmitted in parallel, sequentially to a four-stage, sixty-four bit per stage, buffer. The last stage of this buffer is connected to a series of sixty-four tri-state bus drivers for driving parallel bus 105. Also connected to parallel bus 105, are n output buffers 111, . . . , 112. These outputs buffers also comprise four-stage, sixty-four bit registers, the input stage of which is connected to sixty-four bus receivers connected to bus 105, and the output stage of which is connected to a shift register for generating a serial output stream. Also connected to bus 105, is microprocessor 120 which accepts inputs in bursts of 256 bits as four associated 64-bit data bus reads, from each of the n input buffers 101, . . . , 102, under the control of the program stored in the microprocessor. Similarly, the microprocessor delivers bursts of 256 bits as four associated 64 bit data bus writes to each of the n output buffers after having generated the output burst through the reading of the inputs under the control of a control map, and of the program of the microprocessor.

[0017] An I/O decoder unit 130, under the control of the microprocessor, is used to gate the tri-state outputs of the input buffers onto the bus, and to gate the output of the bus into the n output buffers 111, . . . , 112. The I/O decoder receives inputs from the microprocessor address bus.

[0018] Also connected to bus 105, is a memory 122 for storing infrequently used data and program text such as data required for performing tests or diagnostics, non-cached TSI code, and as a backup for data stored

in the microprocessor cache, such as the microprocessor program text and the path memory. Also connected to bus 105, is control register 124, which interfaces with a call processing controller or other switches of the telecommunications network, and receives and transmits control messages.

[0019] Figure 2 is a block diagram of those key parts of the microprocessor which are pertinent to the understanding of an invention. The microprocessor contains a program cache 201 for storing the control program which controls the operations of the time slot interchange unit. The output of the program cache goes to an instruction queue 203 for storing a plurality of instructions in order to allow for the rapid execution of simple loops that is made possible using pipelining techniques. The instruction queue interacts with an instruction Control Block 205, to deliver the appropriate instructions to arithmetic and logic unit (ALU) 207. The ALU executes its received instructions and operates to perform the steps required by the instruction, by controlling load store unit 213, which in turn accesses a data cache 211. ALU 207 also controls a group of internal registers 215, for short term storage, and for the control of the microprocessor. A bus interface 217 communicates between bus 105, (Fig. 1), and within the microprocessor with the data cache 211, and for changes or back-up in the software, also communicates with program cache 201.

[0020] Figure 3 shows pertinent memory data stored in data cache 211 of microprocessor 120, and in hardware registers. The contents of the data cache contain, among other items, the data received from input buffers 101, . . . , 102, and the data to be delivered to output buffers 111, . . . , 112. Data received from the input buffers 101, . . . , 102, is stored in TSI buffer 301 or 303. The data from the various input buffers is stored sequentially in one of these buffers in Applicants' preferred embodiment. In order to handle nx64 kilobit per second connections, the TSI buffer contains buffer 301, and a second buffer 303, for storing another frame of this serial input data. Buffers 301 and 303 are used alternately. Control map 311 is used to control the reading of the contents of TSI buffers 301 or 303 in order to generate an output for storage in the TSI output buffer 321, for transmission to one of the output buffers 111, . . . , 112. TSI write pointer 315 is used to keep track of where the next input from one of the input buffers 101, . . . , 102, is to be stored in TSI buffer 301 or 303. Control pointer 313 is used to point to the appropriate portions of control map 313 in order to control accessing the TSI buffer in order to obtain the time slots that are required to fill the TSI output buffer 321. Input buffer count 331 is used to control the cycling for accepting inputs from the appropriate one of the n input buffers 101, . . . , 102, selected by input buffer address register 332, and output buffer count 333 is used to control the distribution of an output collected in TSI output buffer 321 to one of the n output buffers 111, . . . , 112, selected by output buffer address register 334. Link status memory 341 is used to identify

any of the n input links or any of the n output links that are out of service. This status can be checked prior to accepting an input from one of the n input buffers 101, . . . , 102, or prior to sending an output to one of the output buffers 111, . . . , 112.

[0021] The control map is altered under the control of the program of the microprocessor when the microprocessor receives a control message from connect request register 351 within the control register 124 of Figure 1, the control message representing a request to establish or disconnect a connection in the time slot interchange unit. The process of controlling the control map is well known in the prior art.

[0022] Figure 4 is a flow diagram describing the operation of the program for implementing a time slot interchange (TSI) in accordance with Applicants' invention. The process starts with the microprocessor waiting for a frame synchronization pulse (Block 401). When the frame synchronization pulse arrives, it signals the beginning of the synchronized loading of the input buffers 101, . . . , 102, from the serial input streams, and triggers several initialization steps. The memory write address, (TSI write pointer 315), is initialized, (Action Block 402), so that the correct location in the TSI buffer 301 and 303 is established for writing the information from the input buffers 101, . . . , 102. The double buffering offset is toggled (Action Block 403), to choose either frame memory 301 or 303 in the TSI buffer for storing the input data on alternate frames. The microprocessor then waits for an input buffer loaded signal, (Action Block 404), which establishes that the buffers 101, . . . , 102, are full, and then the input buffer address is initialized, (Action Block 405), to point to the first input buffer 101. In order to guarantee that Action Block 406 reads new data from the input buffer, and not stale cached data from a previous cycle, Action Block 405 invalidates the cache data associated with the input buffer address before initiating the read. The input buffer pointed to by the input buffer address is then read. (Action Block 406), in a burst as four connected 64 bit data bus operations, and stored in the microprocessor cache memory in either TSI buffer 301 or 303, depending on the double buffering offset. Test 407 determines whether all inputs for this frame have been written. If not, then the buffer address is incremented, (Action Block 409), and the next buffer is read into the TSI buffer, (Action Block 406, previously described). This loop is continued until the results of test 407 indicates that all inputs for this frame have been written.

[0023] At this point, the TSI read cycle begins. The output buffer address 334 is initialized, (Action Block 421), the TSI output buffer address is initialized, (Action Block 423), and the control map pointer 313 is initialized to point to the top of the control map, (Action Block 425). The contents of the control map are read to an index register, (Action Block 427), and the index register is used to read the eight bit time slot from the TSI buffer, (Action Block 429); (frame 301 or 303 is accessed de-

pending on the double buffering offset established in Action Block 403). The read byte is then written into the TSI output buffer in the cache, at the appropriate offset, (TSI output buffers 321), (Action Block 431), determined by which of the 32 bytes is being written. Test 433 is used to determine whether 32 bytes have been written; if not, Action Block 427 is re-entered, and the loop repeats Action Blocks 427, 429, 431. When 32 bytes have been written, as indicated by a positive result of test 433, then 32 bytes are written from the cache, (Action Block 441), by a data cache block flush operation in a burst of four connected 64 bit data bus writes into the output buffer 111, . . . , 112, specified by the output buffer address 334. Test 443 determines whether all outputs have been written. If not, then the TSI output buffer read address is re-initialized, (Action Block 445). The output buffer address, (output buffer address 334), is then incremented, (Action Block 447), and the loop for writing into the output buffer is re-entered in Action Block 427. If test 443 indicates that all outputs have been written, then the work for this frame is finished, and the processor goes back to Block 401 to wait for the next frame synchronization pulse.

[0024] The above flow chart provides double buffering for all time slots whether they represent $n \times 64$ kilobits per second signals such as 256 kilobit data, or a single 64 kilobit per second voice or data time slot. If the additional frame delay introduced by the double buffering is not desired for the single 64 kbit/sec voice or data time slot, then the flow chart can be modified to provide selective double buffering, i.e., the single voice or data time slot is not double buffered. Such single buffered time slots are marked in the control map 311, which causes the time slot to be read from the other one of the two TSI buffer frames 301 and 303, by negating the effect of the double buffer offset. Thus, single buffered time slots may be read out of the opposite frame from the double buffered time slots.

40 Generalized TSI Flow

[0025] The flow chart shown in figure 4 is transversed only once per frame because each of the serial input streams was assumed to consist of 32 time slots, which in the present implementation is written into the microprocessor cache in a single 32 byte burst, as described when Action Block 406 was discussed. Relatively simple modification of Figure 4 is required and illustrated in Figure 5, in order to accommodate higher bandwidth serial links:

(1) Another decision state 451 is required after the "Yes" output of decision state 443 in Figure 4. This determines whether the entire frame of time slots has been processed. If "Yes", we return to the wait state of Block 401. If "No", we return to the wait for input buffer loaded block 404 for the next burst of 32 time slots.

(2) The initialize read control map pointer, Action block 425 is moved out of the TSI read loop to the beginning of the TSI write cycle, (after the initialized memory write address, Action Block 402), since the entire frame has not yet been written.

[0026] The RISC microprocessor hardware of Figure 1, the block diagram of Figure 2, and the programmer data model of Figure 3, can also be used for implementing a Time Multiplexed Switch, (TMS). The basic difference is that a TSI application requires the storing and maintaining in memory of either one or two frames of time slots, (single or double buffered applications), whereas a TMS application requires the switching of the time slots as soon as possible after they appear at the input to the TMS. This means that after the serial input streams appearing at 101, . . . , 102, (which have been written into the TSI buffer of Figure 3), are read out to the serial output streams 111, . . . , 112, their storage in the TSI buffer is no longer necessary. Therefore, subsequent write bursts into this buffer during the frame interval can overwrite the old data. This means that less memory is required for the TMS application than the TSI application, since only 32 bytes, (the write/burst size), per serial input are required rather than one or two frames of memory. Also, double buffering is not required for nx64 kbits per second, because the time slots are read out immediately, and thus, there is no possibility of the time slots getting out of sequence.

[0027] Figure 6 is a flow chart for implementing a TMS. It is similar to the TSI basic flow chart (Figure 4), and incorporates the changes described earlier for generalized TSI flow, as well as the differences described above for a TMS. To help the reader, the same action is given the same number as in Figure 4. For a TMS, high bandwidth facilities, much larger than the 2.048 Mbits per second assumed for the basic TSI flow chart, are required. This requires the addition of test 449 in Figure 6 in order to handle the entire frame, and moving the initialize read control pointer, (Action Block 425), from the TSI read cycle to the frame initialization portion near the beginning of TMS write cycle of Figure 6. These two steps are the same as those described for a generalized TSI flow.

[0028] To implement TMS functionality, the only two changes to the flow chart are:

(1) Move Action Block 402 from the frame initialization portion of the TSI write to the buffer loaded inner loop, so it can overwrite the previous burst, since as described in the previous paragraph, this data has already been output; and

(2) eliminate Action Block 403, which is used to implement double buffering. The TMS flow chart of Figure 6 implements the time multiplexed switching function.

[0029] A variation on writing the input buffers 101, . . . , 102, sequentially into cache, is that instead of taking a 32 byte burst from a single input buffer, 8 bytes from each of four input buffers are written. This has the advantage of reducing the number of bytes of buffering required by input buffers 101, . . . , 102, from 32 bytes to 8 bytes per buffer. Taking 16 bytes from each of two buffers can also be implemented.

[0030] Figure 7 is a flow diagram illustrating the operation of the system when used for switching groups of time slots at a time. This use would be for replacement of a digital cross-connect such as the DACS, (Digital Access And Cross-Connect System), systems manufactured by Lucent Technologies.

[0031] Blocks 461, 463, 465, and 467, replace the functions carried out by Blocks 429, 431, and 433, in Figure 6. In the implementation described in Figure 7, only Blocks 461 and 463 are repeated 8 times. In the final repetition, 465 and 467 are shown, but instead of using the loop, the program is written in-line. Action Block 461 is essentially equivalent to Action Block 429 of Figure 6, and Action Block 463 is essentially equivalent to Action Block 431 of Figure 6; however, instead of having a test of 433, the code is simply repeated 8 times prior to entering Action Block 441.

[0032] The above flow chart described an eight bit time slot where a byte quantity is read and written in Action Blocks 429 and 431. Sixteen and 32 bit time slots can easily be accommodated with a straightforward substitution of half word or full word microprocessor instructions, for the corresponding load and store byte instructions. The time slot width can be further generalized to include group switching, where contiguous time slots are switched as a group using load/store string instructions in Action Blocks 429 and 431, to transfer a sequence of time slots. The total number of bytes of switched information per unit of time, increase with increasing time slot width or group size, since the loop overhead of Action Blocks 427 through 433 is reduced proportionally relative to that of a byte wide time slot. This is very efficient for switching a 32 time slot PCM (E1) facility, for implementing a cross-connect. Some group sizes like that of a T1 facility of 24 byte wide groups, might be most efficiently switched by padding the 24 time slots to a 32 byte group. Groups can be concatenated contiguously to form higher bandwidth rates, such as DS3 at the output of the output buffers; this is especially useful for performing the function of a digital access and cross-connect system.

[0033] The block diagram of Figure 1 can also be used to implement an ATM Switch. Figure 8 shows the structure of an ATM cell header. The ATM cell might be most efficiently switched by padding the 53 time slots into a 64 byte group. This requires some control logic in the Input Buffers and Output Buffers. The generic flow control bits 5 - 8 of octet 1 are used for overall control to prevent an ATM system from being overloaded. The virtual path identifier is split across the first four bits of the

first octet, and the last four bits of the second octet. The virtual path identifier identifies a user. All virtual channels of the same user use the same virtual path identification. The virtual path identifier is the primary identifier used for switching ATM cells within a switch, and for identifying incoming ATM cells so that they can be switched to the appropriate destination. The virtual channel identifier, (the first four bits of octet 2, all of octet 3, and the last four bits of octet 4), is used by the user to identify a specific communication among a plurality of the communications between the end users; the specific communication resides on a specific channel. The first four bits of octet 4 are the payload type (2 bits), one bit reserved for future use, and a cell loss priority bit. The cell loss priority bit is used to help determine whether a particular cell may be discarded in case of overload. Finally, the header error control octet is a cyclic redundancy check (CRC) over the header.

[0034] Figure 9 is a functional overview of the software control components of the ATM Switch. They consist of a CRC Check, Input Link Control, VPI/VCI Processing, Shaping, Quality of Service (QOS) Processing, Output Link Control, and CRC Generate. The cyclic redundancy check, (Action Block 901), is performed on the header of each ATM cell as it enters the system. Input link control, (Action Block 903), brings incoming data into the memory of the microprocessor. VPI/VCI processing, (Action Block 905), finds a VPI/VCI data block which contains an Input VPI/VCI indication, an Output VPI/VCI indication, and a Quality of Service, (QOS), pointer. Test 907 is used to determine whether a shaping test is necessary. Shaping tests are not performed on every cell, but typically, on every tenth cell. If this is a cell that requires the performance of the shaping function, this shaping function is executed, (Action Block 909). The shaping function determines whether the peak or the average allowed data rate is being exceeded. If so, the shaping function introduces a throttle to the transfer of information, which is regulated by putting packets into a shaping queue with limited size, so that if the peak rate is exceeded for too long a time, or the average rate is exceeded, there would no more space in the shaping queue, and the input would be throttled, or packets would be dropped.

[0035] Next, Quality of Service processing, (Action Block 911), is carried out. Each output link has a plurality of queues to provide cells to that output link. The queues contain information of different priority so that certain queues are served preferentially compared to other queues. Finally, the output link control, (Action Block 913), transmits cells from one of the QOS queues to an output link, and a new CRC is generated. Prior to inserting the cell into one of the QOS links, the output VPI/VCI is inserted into the cell header. For some implementations, the CRC functions can be done in hardware in order to increase the switching capacity of the ATM Switch.

[0036] Figure 10 shows the Programmers' Data Mod-

el including register assignments and the data structures used in the implementation. ATM cell routing, defined by the Virtual Path (VP), and Virtual Channel (VC), identifiers is implemented by table look-up in an off chip Static Random Access Memory (SRAM), or in a Level 2 cache, using a hashing algorithm. Queuing of cells is implemented by means of a shared buffer area in cache memory and linked lists associated with each of the output ports. There is also a linked list associated with the unused memory locations, which is used as a pool for adding members/locations to any of the linked lists. Each output link has multiple output queues, each of which is associated with a specific Quality Of Service (QOS). Each output link uses a table look-up of priority, to give the identity of the next QOS queue to be output. This allows the QOS queues to be accessed in any priority sequence desired.

[0037] While in this preferred embodiment, everything is in the cache, for other implementations, especially those with high throughput, much of the data, and some of the more specialized programs, can reside in an external memory.

[0038] The function of the various Blocks of Figure 10 is as follows:

Block 1001 represents the input buffers to the switch.

The input buffer address register, 1003, determines which buffer the system is processing.

Cell header address register 1005, and cell header register 1007, are used for processing the header of one particular cell.

Block 1009 is used for checking and generating the header CRC, (in some alternate configurations, the CRC can be checked or generated automatically by circuitry).

Blocks 1011, the hashing function register, and 1013, the hashing product register, are used for locating the VPI/VCI specified in the header of an input cell.

Block 1015 is the VPI/VCI Table, which is typically occupied only 50 percent to allow for efficient hashed access.

Some of the blocks pointed to by Table 1015, are Block 1017, which is the VPI/VCI block for VPI/VCI 1, Blocks 1019, which are empty blocks, and Block 1023, the block for the last VPI/VCI.

Block 1017 includes the identity of the input VPI/VCI, the identity of the output VPI/VCI to which the cell should be switched, and a pointer to the Quality of Service, (QOS), queue which is used for assem-

bling cells to be transmitted to the output link.

[0039] The third column of Figure 10 shows a plurality of QOS queues, using a shared memory spectrum, one set 1031, . . . , 1033 for link 1, and another set, 1035, . . . , 1037, for the last link, Link "n". Block 1031 includes an identification of the link for which cells are being queued, and a pair of pointers for the entries in the queue. The entries in the queue are linked each to the next, and the head cell pointer is used to find the cell in the queue which is to be transmitted to the output link, and the tail cell pointer finds the entry in the queue in which the next cell can be entered. Finally, Blocks 1041 and 1043 are used for selecting the particular cell in one of the QOS queues which is to be transmitted to an output buffer. For each output buffer, there is one link control, such as link control 1043. Link control 1043 contains head cell pointers to the QOS queues. For high priority QOS queues, several entries would be made in the table of 1043, which has sixteen entries, with the sixteen entries being more than the typically 4, QOS queues per output buffer. The output link register is used to select which link is being processed, and the priority counter register is used to select the head cell pointer for that output buffer. When the head cell pointer of Block 1043 is read, it will point to a head cell pointer of one of the QOS queues, and that head cell pointer in turn, will point to the oldest cell in that queue, i.e., the cell which is to be placed in the output buffer. Finally, Block 1051 shows the "n" output buffers, output buffer 1, 1053, . . . , output buffer "n", 1055. The output address register 1057 is used to select which output buffer is being processed.

[0040] Figure 11 is a flow chart showing the Cell Input and VPI/VCI Flow. The cell input section shows the writing of a burst of 32 bytes into the cache memory from the Input Buffer selected by the Input Buffer Address. The header and VPI/VCI Processing are shown in the remaining part of the figure. The CRC check can be done in software if desired, and is implemented by using the header, a byte at a time, for indexing into a 256 byte table. If an error is detected, a routine is entered which either corrects the error or results in the cell being dropped. After the CRC check, an Empty Cell Code Check is done. Empty cells are ignored, but the routine goes to the normal "single thread" output routine, ("E" input of Figure 13). Next, a 32 bit hashing function is used in conjunction with the VPI/VCI, to generate a hashing address for indexing into SRAM or Level 2 cache, and read a 32 byte burst of data for that VPI/VCI. If the correct VPI/VCI is not at that address, alternative hashing addresses are iteratively tried until either the correct VPI/VCI is found, or the exception handling routine, is entered. Hashing algorithms are well described in the literature. For a VPI/VCI table which is only 50% occupied, the average number of searches required by the implemented algorithm is 1.5, thus, providing reasonable access times at the expense of memory. When the search is successfully completed, shaping is per-

formed, if necessary, and the "Output VPI/VCI", i.e., the destination for the cell, is extracted from the table and inserted into the cell header.

[0041] Figure 12 is the Output Queue Flow Chart. It consists of inserting the cell into the appropriate output queue based on the output link, and the QOS specified in the data associated with the VPI/VCI search described in the previous paragraph. There are "m" QOS queues associated with each output link, and each queue is defined by a linked list, (see the "m" QOS Queues Per Output Link Tables in Figure 10). Linked lists are well-known in the prior art. There is also a list of all the unused memory locations defined by an unused locations link list, called an "Unused Location Queue", (ULQ). Figure 12 details the pointer and data manipulation to implement the linked list queues.

[0042] Figure 13 is the Write to Output Links flow chart. The priority sequence used for the output queues is to use a static Per Output Link Priority Table, (see Figure 10), to establish the sequence of queue readout on a per link basis. The Per Output Link Priority Tables shown in Figure 10, show, (as an example), 16 entries, each of which could specify any of the "m", (e.g., m=4), queues established for that link. If the selected queue on a link is empty, each of the other queues are interrogated until a queue with data is found, or it is determined that all of the queues associated with the link are empty. If a cell is present in any of the queues, then the CRC is generated and inserted in the header and the cell is transferred to the output buffer. If there is no cell in any of the queues, then the CRC for Idle Code is generated, and an Idle Code cell is transferred to the output buffer. There is then some pointer manipulation associated with housekeeping of the linked lists. There is further housekeeping associated with priority and buffer address manipulation. There are also some decision points regarding All Links Written, Shaping, and All Cells Read resulting in appropriate loop back to entry points in Figure 11, or transfer to the shaping routine.

[0043] Shaping, (Action Block 909), occurs at multiple, periodic cell intervals to assure that the per VPI/VCI contracted peak and average bandwidths are not being exceeded. Cells can either be dropped, delayed or passed through. Shaping is done on a per VPI/VCI basis using linked list auxiliary queues. The details for performing shaping are well known in the prior art. Additional information is stored in the VPI/VCI Table of Figure 9. For the shaping interval being considered, (e.g., every 10 cells for peak rate, and every 100 cells for sustained rate), the following information is provided in the VPI/VCI table: contracted Peak Cell Rate, (PCR), time stamp for PCR, contracted Sustained Cell Rate, (SCR), time stamp for SCR, and maximum size of shaping queue.

[0044] The individual Blocks of Figures 11 - 13 will now be described. Figure 11 starts in Block 1101, wherein the system is waiting for a frame synchronization pulse. When the frame synchronization pulse arrives, it

signals the beginning of the synchronized loading of the input buffers 101, . . . , 102, (Figure 1), in serial input streams. Action Block 1103 indicates a wait for the signal that input buffer has been loaded. When the input buffer has been loaded, the memory write address for unloading that buffer into microprocessor memory is initialized. The cell is then read from the input buffer, (Action Block 1107), and the input buffer address is incremented, (Action Block 1109). At this point, the cell has been loaded into the memory of the microprocessor, and the microprocessor is ready to process the cell. The header of the cell is loaded into a register, (Action Block 1121), and a CRC check is performed, (Action Block 1123). A CRC check is performed only on the contents of the header. A CRC check can be performed with special circuitry, or it can be performed relatively expeditiously through the use of a table of 256 bytes; each byte corresponding to one of the 256 possible CRC bytes. Next, a check is made to see if the cell is empty, (Test 1125). An empty cell has an industry standard predetermined VPI/VCI identification. Test 1127 determines whether the cell is in fact empty, and if so, further processing is terminated, and the output processing routine of Figure 13 is entered. If the cell is not empty, then the VPI/VCI Table entry for this cell must be found. Action Blocks and Tests 1129, 1131, 1133, 1135, 1137, 1139, 1141, 1143, and 1145 describe this process. The VPI/VCI Table, (Table 1015), of Figure 10 is found, (Action Block 1129). A hashing function, a known constant, is then loaded into a register of the microprocessor, (Action Block 1131). This register is then multiplied by the contents of a register containing the VPI/VCI, (Action Block 1133). In one example of this embodiment, there are up to approximately 2,000 VPI/VCI entries, such as Block 1017 of Figure 10. In the Table, 12 bits of the product generated in Action Block 1133, the least significant 12 bits in this case, are then used to read an entry in the VPI/VCI Table. The Table is 4,096 entries long, and corresponds to the 12 bit accessing queue. In Action Block 1137, the actual VPI/VCI is compared with the VPI/VCI found in the accessed VPI/VCI Table, (Action Block 1137), if Test 1139 is used to determine if the two are equal. Equality means that the appropriate VPI/VCI Table entry has been found. If not, then Test 1141 is used to determine whether this is already the "nth try", and if so, the exception handling routine 1143 is entered. This routine searches a list of VPI/VCI Table entries, (not shown in an Auxiliary Table), used for serving cases in which "n" tries fail to locate a VPI/VCI. Entries in the Table are created in those cases where an attempt to load the Table, encounters "n" failures. If this is not the "nth try", then a different 12 bits of a 32-bit product generated in Action Block 1133 is used, (Action Block 1145), in order to access a different entry of the VPI/VCI Table, (Action Block 1135).

[0045] The hashing arrangement is used because the total number of possible VPI/VCI combinations is over a million, (the VPI indicator is 8 bits long, and the VCI

indicator is 12 bits long), so that 2^{20} (more than one million), possible values of VPI/VCI exist even though only 2,000 are being used at any one time.

[0046] Once the appropriate VPI/VCI Table entry has been found, (with output of Test 1139), Test 1151 is used to determine whether shaping is required in this case. In this embodiment, shaping actions are performed only on every "nth" cell, wherein "n" may, for example, have a value of 10. Shaping is used to monitor the input rate of a particular VPI, to insure that the VPI does not send more cells than is allowed for its peak rate. The peak rate is defined as the number of cells which may be sent for a particular interval. If more than this number of cells is sent, then the extra cells are either simply discarded, or are temporarily passed on, but a slow-down message is sent to the source of the cells. After the shaping function has been performed, (Action Block 1153), or in case shaping is not required for this cell, then the output VPI/VCI identifier is loaded from the VPI/VCI Table into the cell, and is substituted for the input VPI/VCI. Thereafter, the output queue routine of Figure 12 is entered.

[0047] The system reads the QOS pointer stored in the VPI/VCI Block. This pointer points to a tail cell pointer within the QOS queue for serving that VPI/VCI. The QOS queue, (for example, Block 1037 of Figure 10), is used to queue cells for transmission to an output link. As previously mentioned, several QOS queues serve a particular output link, and depending on the quality of service being supplied to a particular VPI/VCI, the cells are stored in a different queue, and different QOS queues are served preferentially, for delivering their contents to an output link. The contents within each QOS queue are stored in a linked fashion, and the last entry is pointed to by a tail cell pointer. It is this pointer which is pointed to by the QOS pointer in the VPI/VCI Block. The QOS queue "m" pointer is read, (Action Block 1203), and a "n" link from that idle queue location to the next idle queue location, is temporarily stored in a register of the microprocessor, (Action Block 1205). The cell is then stored in the queue at the location originally specified by the "n" cell pointer, (Action Block 1207), and the address of the next empty cell.

[0048] In order to share the available memory space effectively and dynamically, linked lists are used for each of the output queues. In addition, there is an "unused location" linked list, which is a global resource containing the empty, (unused), locations available for storing information in queues. When a queue needs to add information, it gets the available location from the "unused location" linked list. As a result, both the "unused location" linked list, and the linked list of the queue requesting an available location, are impacted. There is a separate head cell pointer, and tail cell pointer associated with every queue, including the Unused Location Queue, (ULQ).

[0049] The head cell of the ULQ is the next available location for storing a queued cell, and the tail cell of the ULQ is the last cell that has been returned to the ULQ

pool. The head cell of a queue is the last cell that has been stored in that queue, and the tail cell of a queue is the next cell to be out-putted from that queue. The head cell of the ULQ becomes the tail cell of the queue that is requesting a storage location, and the linked lists of both are modified to support this transfer of memory location function. Specifically, Action Block 1204 extends the queue to include the cell stored by Action Block 1207, and Action Block 1211 updates the pointer to reflect this linked list extension. Action Block 1213 changes the head cell pointer of the ULQ to reflect the removal of an available cell location.

[0050] Following the execution of Action Block 1213, the output processing of Figure 13 is performed. Block 1043 of Figure 10, is a series of 16 pointers to the "m" QOS queues of a particular output link, where "m", in this example, is much less than 16, typically 4, so that the 16 entries can be used to service different QOS queues more, or less, frequently. Associated with an output queue is a priority counter 1045, which is used to select the appropriate entry from the priority table. In Action Block 1301, the priority counter is used to index into the priority table of the output link being serviced, (different output links are serviced on a rotating schedule). The priority counter is then incremented in order to prepare for servicing the link the next time, (Action Block 1303). The queue pointed to by the priority table, is then checked to see if it is empty, (Action Block 1305). Test 1307 is used to determine whether the queue is empty, and if so, whether this is the last, (4th), queue, (Action Block 1309). If it is not, then the queue counter is decremented, (Action Block 1311), and the corresponding queue is checked to see if it is empty, (Action Block 1305). If the result of Test 1307, either initially, or after having gone through the loop, using 1309, 1311, and 1305, indicates that the queue is not empty, then a CRC is generated for the cell header, (Action Block 1313), and the cell header is stored in the output buffer, (Action Block 1315). The output buffer address is incremented to prepare for subsequent processing, (Action Block 1317), and the queue this cell was transferred to the output buffer, is updated to add the storage of the cell which was transferred to the buffer to the list of empty locations in the queue, and to update the head cell for the queue.

[0051] Action Blocks 1321 to 1325 represent link list pointer manipulation for reading from a queue to the output link, and is similar to the write sequence for Action Blocks 1204, 1211, and 1213, described above. In this case, however, a cell location is added to the ULQ pool, and a cell location is removed from the queue that has out-putted a cell.

[0052] Test 1335 is then used to determine whether outputs to all links have been sent. If not, the output link priority table is incremented, (Action Block 1337), so that at the next pass, the next link will be served. Action Blocks 1339 and 1341 are used to unload the shaping queue. In the case that outputs to all links have been generated, (positive result of Test 1335), then the output

link priority counter is incremented, 1351, the input buffer address is initialized, (Action Block 1353), so that the first input buffer is then serviced, the output buffer address is initialized, (Action Block 1355), so that at the next pass, the initial output buffer will be serviced, and the output link address register is initialized, (Action Block 1357). Test 1359 then determines whether all cell have been read from the input buffers, and if not, Action Block 1103 of Figure 11, is re-entered. If all cells have been read, then Action Block 1101 of Figure 11 is entered.

[0053] Figures 11 to 13 show the flow charts for implementing an ATM switch, exclusive of the shaping, (which occurs only at multiple cell intervals to assure that the per VPI/VCI contracted peak and average bandwidths are not being exceeded). The flow chart deliberately shows a "single thread" implementation, i.e., one cell at a time is taken from input to output before the next cell is input in order to demonstrate the logic of the design. Efficiencies in processor utilization can be obtained by overlapping functions such as I/O read/writes, and read/writes of off chip memory and Level 2 caches, by doing "multiple thread" ATM cell processing.

[0054] The above implementation of ATM switching assumed that the ATM cells coming into the switch were in the format of 53 contiguous time slots, which characterizes an important segment of the applications. There are other applications where an ATM cell comes in over lower bandwidth pipes, e.g., fractional T1/E1 using 128 Kbps, 384 Kbps, etc. For those cases, the ATM cell needs to be aggregated by examining a number of frames until the entire 53 byte cell is available. There are several ways to implement this. One way is to consider this function as part of the periphery, and provide a separate RISC microprocessor to provide the function. A second way, is to incorporate the aggregation function into the ATM switching fabric discussed above. Different tradeoffs will exist for different applications, e.g., the ratio of fractional ATM to complete cell ATM, as well of the size of the switching job being considered and the amount of real time available.

[0055] The block diagram of Figure 1 can be used to implement an Internet Protocol, (IP) switch, as well as an ATM switch, whose functionality is described in Figs. 9 - 13. Unlike the ATM case, an IP packet is of variable length, and has a destination address field that requires a longest prefix match for switching. The variable length implies more flexible buffer allocation schemes, and potentially, requires packet fragmentation and reassembly, depending upon the maximum transmission unit sizes in the different networks that the IP switch would switch between. The sequence of processing steps can be similar to the ATM case, and would consist of header checksum verification, input link control, destination processing, quality of service processing, output link control, and header checksum verification. In some implementations, the header checksum processing could be done in hardware in order to improve the capacity of the IP

switch.

[0056] After IP header checksum verification, IP packet routing examines the destination address field of the IP header, and performs a hash based look-up algorithm that can search for a longest prefix match as is well described in the literature. The search would return information about the appropriate output link. Further analysis of the packet header could yield treatment information for implementing various levels of quality of service, and would locate a specific output queue associated with the output link and the assigned treatment quality. If the packet length is larger than the maximum transmission unit size of the output link, then the packet would be fragmented and linked to the appropriate output queue as a sequence of packets. Output link processing would select a packet from the highest priority queue of the moment, and perform final adjustments to the selected IP header such as adjusting the time to live field and the header checksum of the modified IP header before committing the packet to the actual physical output link. A time to live field is used to discard an Internet packet if it is not delivered within the time, or the number of switching points traversed, specified in the field.

[0057] IP, (Internet Protocol), switching can be performed within an universal switching fabric via software emulation of functionality, that would in less flexible implementations, be performed in hardware, often in Field Programmable Gate Array, (FPGA) based state machines. In all implementations, well formed packets are eventually handed to switching and routing software. The headers of these packets would be examined for classification as to flow types via hashing to determine output queues. The flow classification could use various protocol and port data from the packet to be switched in addition to the IP destination, in forming keys to the hashing process. The hashing search ultimately yields output link and queue information allowing for, (Quality of Service), QOS treatment. Various IP fields such as (Time to Live), TTL, would be updated as the packet was linked to output queuing. The routing information embodied in the flow based hashed search table, would be maintained through gateway protocol processing. Output handling would on a per link basis, always determine the next best output queue, to unlink a packet for actual packet transmission. As is described in the above packet formation case, the packet output case could also be embodied in several different implementations. IP switching uses many of the mechanisms described in more detail in the ATM section. Depending upon performance trade-offs, different embodiments of these concepts can move the functionality of packet formation from serial streams; various separate sequentially cooperating processors can be used instead of a single processor to form packets from within TSI time-slot locations marked as containing packet stream data.

[0058] Frame relay switching can also be implemented within a software based universal switching fabric. In the frame relay case, HDLC based processing would be

best performed by input adaption hardware, because the bit oriented processing would often not be cost effective in universal switch software. Assuming that well formed frames were handed to the frame switching software, hash searching over DLCI field information would yield output link and queue information. Separate Operations, Administration and Maintenance, (OA&M) software would maintain the frame routing information embodied in the frame hash route table. Subsequent output processing would unlink the frame from its output queue for transmission within an HDLC format by output adaption hardware.

[0059] So far, this document has described examples of single function switching fabric implementations which can reside on, and be implemented by a common RISC microprocessor architecture. These single function switching fabrics can reside, and be implemented concurrently on the same microprocessor.

[0060] In its simplest form, the different type of fabric functionality can be allocated on a per serial link interface to the shift registers shown in Figure 1. This would be done under the control of software that can be downloaded, as required. For each type of link, the program for processing the protocol of that link is executed when processing that link. For example, if ATM time-slots and circuit switched time-slots destined for TSI functionality occupied separate serial link interfaces, the link time-slots would be burst into the cache as described in the single function implementations. The bandwidth of these serial links, e.g., number of time-slots, could vary depending on the application and the specific serial link. Since TSI time-slots must be retained for one or two frame intervals, (depending on whether the time-slots are single buffered or double buffered), reading in of subsequent ATM cells which do not have this frame retention requirement, could result in the corruption of the TSI data in the cache. If cache lines are locked after each input burst until the data is no longer required, then this potential problem is avoided.

[0061] This can be extended to more than two concurrent fabric types, including positional switching, (e.g., TSI, TMS, and XCON), and packet switching, (e.g., ATM, IP routing, and Frame Relay). Allocation to individual serial links may be unnecessarily restrictive, for many applications and the different types of traffic can reside on the same serial link with specific chunks of bandwidth being allocated for each protocol type, for switching data being transmitted in different protocols. This could also be done by downloading the appropriate data or software. This could be done using a "recent change" mechanism, as customers select or change their service type.

[0062] In the descriptions for TSI interfaces, the description indicated that the 24 bytes could be padded to 32 bytes in the input/output shift registers. Similarly, for ATM interfaces, it was suggested that the 53 byte cell could be padded to 64 bytes in the input/output shift registers. Although this is reasonable to do when only a sin-

gle functional type is allocated to a shift register. it may add too much complexity when multiple functional types are allocated to a specific shift register. Thus, it may be preferable to read, (burst in), or write, (burst out), the links to/from cache as they are, i.e., a contiguous stream of time slots, and do the padding manipulation in software, inside the microprocessor.

[0063] Bandwidth allocated for different traffic types within a given serial link could be flexibly manipulated by the subject microprocessor, with a linked list of data structures being used to describe sequential memory bytes from each serial interface. Separate input and output lists for each interface could be interpreted by the microprocessor with descriptive codes indicating the traffic type, with length information and application specific points and indicators that could, for instance, indicate where circuit switched data belong within a TSI, or where packet data would be buffered, for reassembly. The switch retains control data for a frame.

[0064] For example, the microprocessor might interpret from the linked data structures for a given interface, that the next M bytes of data should be treated as circuit switched data to be sent to the next M sequential locations of a TSI. The next linked data structure might then contain a code and length indicating that the next N bytes contain a part of an IP packet that is being assembled at the reassembly area, pointed to from the data structure. Finally, for example, the last linked data structure might indicate that the next P sequential bytes contain ATM cells.

[0065] Such a linked list of input and output descriptor data structures could flexibly describe any variety of traffic types within input and output interfaces. The descriptors could also indicate how data should be interpreted within distinct virtual tributaries of the same physical interface. OA&M software would be used to maintain the contents of the descriptor data structures.

[0066] Advantageously, these concepts can be used, for example, within a single microprocessor universal switch application at a small business, where in the prior art, the small business would lease separate fractional T1 facilities, with one T1 facility for PCM circuit switched voice traffic, another for frame relay based IP world wide web traffic, and yet a third T1 facility, for ATM based video conferencing. The leased cost of these separate facilities would often be substantially more than the cost of a single facility, even when more bandwidth would be available if a universal switching element could be used. The universal switching element can also, advantageously, offer the dynamic adjustment of bandwidth between different traffic types within the consolidated leased T1 facility.

[0067] The result of putting multiple concurrently running switching fabrics in a single microprocessor, is to have a minor impact on the capacity of the switches because of the real time impacts described above. It is estimated that a 300 MHz EC603e PowerPC can support about 480 Mb/sec, (7500 time-slots), of single time-slot

TSI switching, or about 1.5 Gbit of ATM cell switching, (3 million cells per second). When sharing functionality on a single microprocessor, the capacity of each of the applications is reduced by the proportional amount of their real time usage. For example, a single microprocessor could concurrently support about 240 Mb/sec, (3750 time-slots), of TSI, and 750 Mb/sec of ATM cell switching, (750,000 cells per second). The ratios for a particular application would depend on the traffic mix, and could include proportional amounts of Frame Relay and IP router switching.

[0068] The above sections have demonstrated concurrent operation of circuit switching and packet switching fabrics. The RISC can also provide the SAC, (Synchronous to Asynchronous Conversion), function required to go between the circuit (synchronous), and the packet, (asynchronous) worlds such as AAL1, AAL2, and AAL5, as well as the layering of IP over ATM, and IP over frame relay. Thus, not only is there connectivity within the each of the switching domains, but also integrated interconnectivity between these switching domains.

[0069] Figure 14 illustrates an arrangement for increasing the size of the TSI of Figure 1. Figure 14 shows an implementation that can be applied to any number n of input signals, any number k of microprocessor complexes, and any number n/k, that can be accommodated by the speed and memory capacity of these complexes. In the specific embodiment of Figure 14, n is 32, k is 8, and n/k is 4. Each of the input streams terminated at the buffer amplifiers 521-1, . . . , 521-32, is connected to a shift register input buffer similar to the input buffer 101. For microprocessor complex 501-1, shift registers 511-1, . . . , 511-32, are connected to local bus 541-1, from which microprocessor complex 501-1 accept inputs. The same arrangement is available for each of the 7 other microprocessor complexes 501-2, . . . , 501-8. Each microprocessor complex feeds only four of the total 32 output buffers. For example, microprocessor complex 501, feeds output buffers 531-1, . . . , 531-4. The capacity of each microprocessor complex must be adequate to take inputs from the full range of input shift registers, but need only drive 1 over k of the output streams. Fortunately, the absorption of the inputs is done in parallel, since input signals are loaded into sequential locations in the TSI buffer 301, 303, of each microprocessor. Thus, very large amounts of input data can be absorbed per unit time in the microprocessor caches. It is only the output data which requires sequential time slot by time slot, or group by group, processing by the microprocessor.

[0070] The arrangement of local shift registers per microprocessor complex has the advantage of limiting high bandwidth connections to the nearby locality of each microprocessor, with the corresponding disadvantage of requiring replicated shift registers for each microprocessor. In another arrangement that might sometimes be advantageous, a single global set of shift reg-

isters could be used, with each microprocessor in lock-step, absorbing the same input data at the same time. In this case, the complexity of high bandwidth global connections and global microprocessor synchronization, would be traded against the savings of a set of shift registers for all but one of the microprocessors.

[0071] Theoretically, it is possible to take input data and process the input data serially in order to generate pre-ordered output data. The arrangement of Figure 14 does not work satisfactorily for that kind of arrangement, (processing inputs serially to generate parallel outputs), because for each input word that is received in parallel, different microprocessors are required to do different amounts of processing, since each processor may process a different number of bytes to generate output streams for its outputs.

[0072] Although the discussion thus far has focused on "switching fabrics", the RISC microprocessor can be used to implement other functionality.

[0073] A RISC microprocessor can also be used to terminate serial telecommunications links. As an example, a link such as a proprietary PCT, (Peripheral Control and Timing) link can be considered. This serial fiber optic link has a 1024 time slots, 768 of which are used for data transport, and the remainder for control, synchronization, and other functions. Frame synchronization is established by a fixed code established in several contiguous time slots. Just as for the switching fabric, the serial bit stream is shifted into an external register, then burst into the cache as bytes of information. The RISC examines the contiguous byte sequence to see if it corresponds to the synchronization code. If not, then a single bit shift instruction is implemented and the resulting, changed contiguous bytes are examined. This procedure of examining the input bit stream for the correct code as subsequent bytes are entered via the I/O, continues until the synchronization code found. This establishes the frame synchronization point, and puts the serial link into synchronization. An additional byte in the sequence establishes the super-frame boundary, which is searched for until super-frame synchronization is also attained. Other required functions can similarly be implemented by appropriate operations on the bit stream. Multiple links can be supported on a single microprocessor.

[0074] Other serial telecommunications links can also be implemented including the well-known standard DS1, DS3, DS_n, E1, E3, and other 32 channel based facilities, SONET, SDH, as well as proprietary serial links such as the PCT, NCT (Network Control and Timing), PIDB (Peripheral Interface Data Bus), etc., used by the 5ESS® Switch.

[0075] One protocol can be used for transmitting data in another protocol. For example, the frame relay protocol, or the ATM protocol can be used for transmitting data in the IP protocol. The switching system can then switch data in the carrying protocol, and the carried protocol data can then be extracted from the switched data.

[0076] A specific microprocessor can be used to terminate either one type of the serial links or several types of these serial links concurrently. The multiple microprocessor configurations shown in Figure 14 can also be used.

[0077] The above has described a Universal Switching Fabric and a Universal Serial Link Termination as separate entities, but they can be combined in a single microprocessor. For example, a single microprocessor can terminate any of the above serial links described, and then provide the ATM switching in the same microprocessor concurrently. If desired, universal serial link termination, (any/all of the above links described in this document), can be coupled with the universal switching fabric, (any/all of the above fabrics discussed in this document), and operate concurrently in the same microprocessor. The multiple microprocessor configurations shown in Figure 14 also apply.

[0078] The above approach can be used to provide higher level combined switching functionality. An example is the implementation of the functionality of an entity such as a trunk only 5ESS® Switching Module on a single microprocessor. This would include the TSI, the trunk terminations, the NCT interface to the TMS, service circuits such as tone generation implemented in a floating point decimal unit of the microprocessor, or vector manipulation unit of the microprocessor as the hardware base. On the same microprocessor, the Switching Module Processor, (SMP), can be concurrently implemented, (either in native or emulation mode), for call processing and maintenance software, as the software base. For an SM (switching module), that contains subscriber lines as well, all of the above can be used to support subscriber lines implemented in the conventional way, as well as the embedded trunk circuits.

[0079] The microprocessor can also be used to implement generalized logic function cost effectively, especially if they have a strong component of sequential logic. Thus, this approach can be useful in providing the functionality presently implemented by a Field Programmable Gate Array (FPGA), and be more cost effective and provide more rapid deployment. This approach can also be used for replacing Application Specific Integrated Circuits (ASICs), with either a single or multiple microprocessor, depending on the application.

[0080] While the preferred embodiment shows sequential storage of input time slots and readout based on the control memory contents, it is also possible to use storage based on control memory contents in conjunction with sequential readout, although such an arrangement handles broadcast connections less efficiently. The arrangement of Figure 14 does not work satisfactorily for broadcasting in the non-preferred arrangement, (storage based on control memory and sequential readout), because for each input word that is received, different microprocessor may be required to do different amounts of processing.

[0081] RISC microprocessor technology is moving at

a very fast pace. There are microprocessors operating at higher frequencies just beyond the horizon which will allow for higher capacities on a single chip. Moore's law will push these capabilities even further in the future.

[0082] The approach discussed herein has the inherent flexibility of stored program control. It can, therefore, be used to implement new and different protocols such as the Dynamic Synchronous Transfer Mode (DTM), recently proposed by the European standards body.

Claims

1. A telecommunications switching network fabric element comprising:

a microprocessor comprising an internal memory;

a plurality of input buffers, each for receiving an input bit stream;

a plurality of output buffers, each for transmitting an output bit stream;

the microprocessor connected to the input buffers by a multi-byte bus for receiving a plurality of bytes; and

the microprocessor operative under a control program for performing switching functions;

wherein different ones of said input bit streams comprise data transmitted in different protocols.

2. A telecommunications switching network fabric element comprising:

a microprocessor comprising an internal memory;

a plurality of input buffers, each for receiving an input bit stream;

a plurality of output buffers, each for transmitting an output bit stream;

the microprocessor connected to the input buffers by a multi-byte bus for receiving a plurality of bytes; and

the microprocessor operative under a control program for performing switching functions;

wherein different ones of said output bit streams comprise data transmitted in different protocols.

3. The telecommunications switching network fabric element of Claims 1 or 2, wherein said microprocessor further performs protocol conversion between protocols of said input streams and protocols of said output streams.

4. The telecommunications switching network fabric element of Claims 1 or 2, wherein one of the input stream protocols is Pulse Code Modulation (PCM).

5. The telecommunications switching network fabric element of Claims 1 or 2, wherein one of the input stream protocols is an Internet Protocol (IP).

6. The telecommunications switching network fabric element of Claims 1 or 2, wherein one of the input stream protocols is an Asynchronous Transfer Mode (ATM) protocol.

7. The telecommunications switching network fabric element of Claims 1 or 2, further comprising:
a plurality of additional telecommunication switching fabric elements, each of which receiving inputs from a common plurality of input streams, and transmitting outputs to a separate sub-set of the output streams.

8. The telecommunications switching network fabric element of Claims 1 or 2, wherein one of the input stream protocol is a frame relay protocol.

9. The telecommunications switching network fabric element of Claims 1 or 2, wherein said microprocessor is further operative under program control for controlling transmission of groups of bytes to each of a plurality of output streams;
wherein channel groups can be efficiently switched as entities for performing a digital access and cross connect function.

10. The telecommunications switching network fabric element of Claims 1 or 2, wherein said internal memory comprises a cache memory.

11. The telecommunications switching network fabric element of Claims 1 or 2, further comprising an external memory accessible by said microprocessor.

12. The telecommunications switching network fabric element of Claim 11, wherein said external memory comprises a cache memory.

13. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one of the plurality of input streams is an input stream carrying data transmitted in more than one protocol.

14. The telecommunications switching network fabric

- element of Claims 1 or 2, wherein at least one of the plurality of output streams is an output stream carrying data transmitted in more than one protocol.
15. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one of the input streams is connected to a single output stream, and wherein the at least one input stream and the corresponding single output stream transmit data in different protocols. 5
 16. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one input stream transmits Internet Protocol, (IP) data carried on an Asynchronous Transfer Mode, (ATM) protocol bit stream. 10
 17. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one output stream transmits Internet Protocol, (IP) data carried on an Asynchronous Transfer Mode (ATM) protocol bit stream. 15
 18. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one input stream transmits Internet Protocol, (IP) data carried on a Frame Relay, (FR) protocol bit stream. 20
 19. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one output stream transmits Internet Protocol, (IP) data carried on a Frame Relay, (FR) protocol bit stream. 25
 20. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one of said input bit streams is a serial telecommunications link, and the microprocessor converts inputs into byte formatted synchronized signals. 30
 21. The telecommunications switching network fabric element of Claim 20, wherein said element further generates output signals in a serial telecommunications link format. 35
 22. The telecommunications switching network fabric element of Claims 1 or 2, wherein said microprocessor further comprises means for executing switching control software; 40
 wherein a single microprocessor controls and executes essentially all switching and switching control functions of a switch. 45
 23. The telecommunications switching network fabric element of Claims 1 or 2, wherein said microprocessor performs ATM header translations, performed in the prior art by an Application Specific Integrated Circuit (ASIC), or a Field Programmable Gate Array, (FPGA). 50
 24. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one of said input streams comprises synchronous signals which are converted to an asynchronous protocol for transmission over one of said output streams. 55
 25. The telecommunications switching network fabric element of Claims 1 or 2, wherein at least one of said input streams comprises asynchronous signals which are converted to a synchronous protocol for transmission over one of said output streams.
 26. The telecommunications switching network fabric element of Claims 1 or 2, wherein said input streams is a frame synchronized stream, and the microprocessor adds additional information required by facility protocols.
 27. The telecommunications switching network fabric element of Claims 1 or 2, wherein said facility protocol is SONET.
 28. The telecommunications switching network fabric element of Claims 1 or 2, wherein said facility protocol is SDH.
 29. In a switching system requiring performance of highly repetitive data processing functions, a method of performing said highly repetitive data processing functions comprising the steps of:
 substituting microprocessor control for Application Specific Integrated Circuits, (ASIC), control.
 30. In a switching system requiring performance of highly repetitive data processing functions, a method of performing said highly repetitive data processing functions comprising the steps of:
 substituting microprocessor control for Field Programmable Gate Array, (FPGA) control.

FIG. 1




FIG. 2



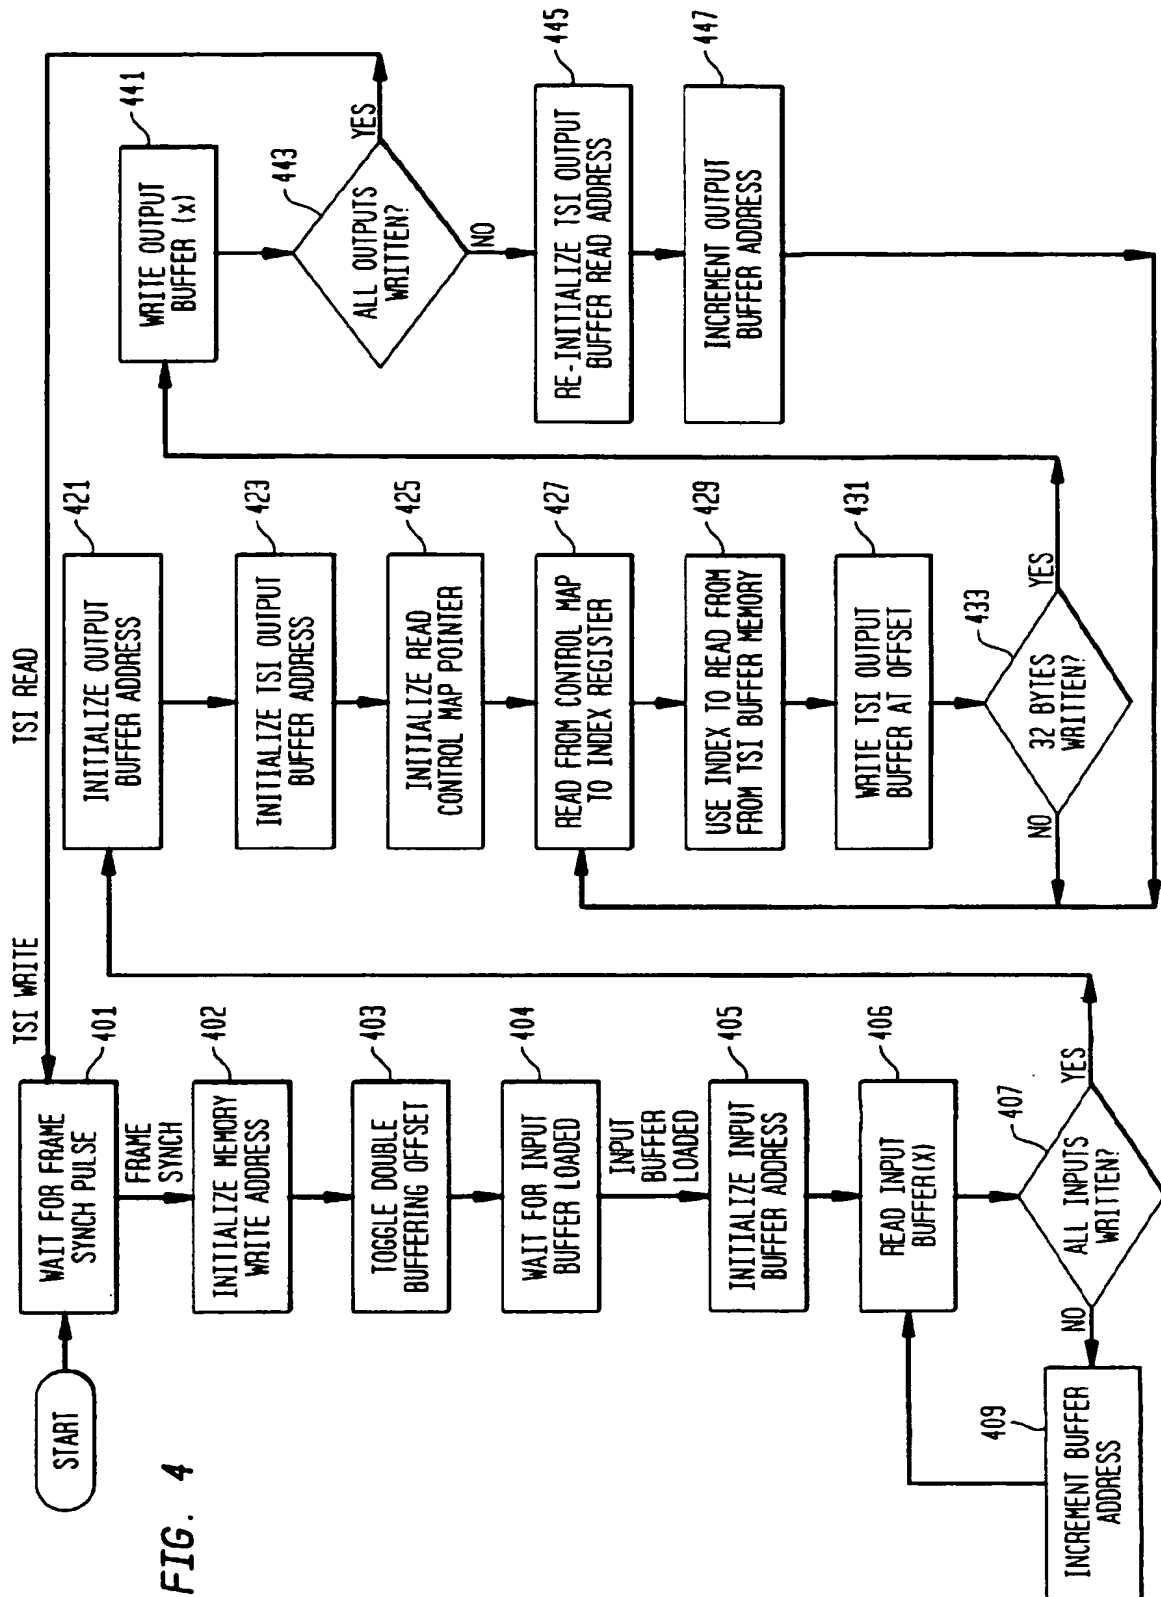
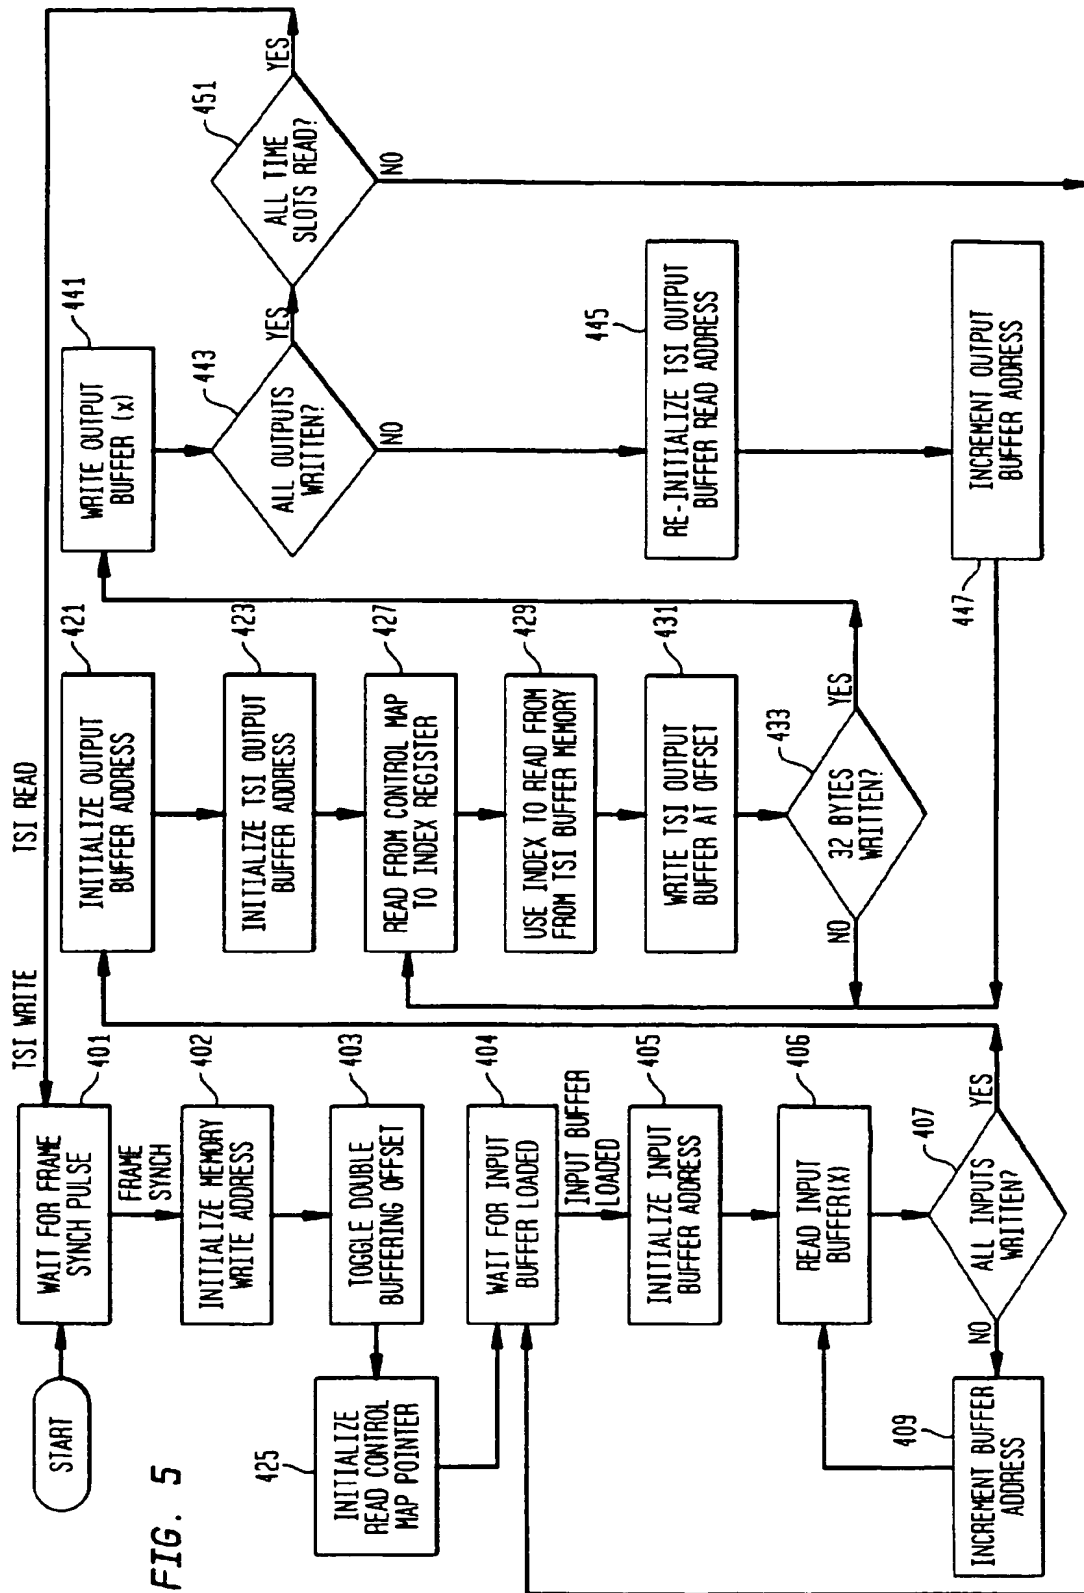
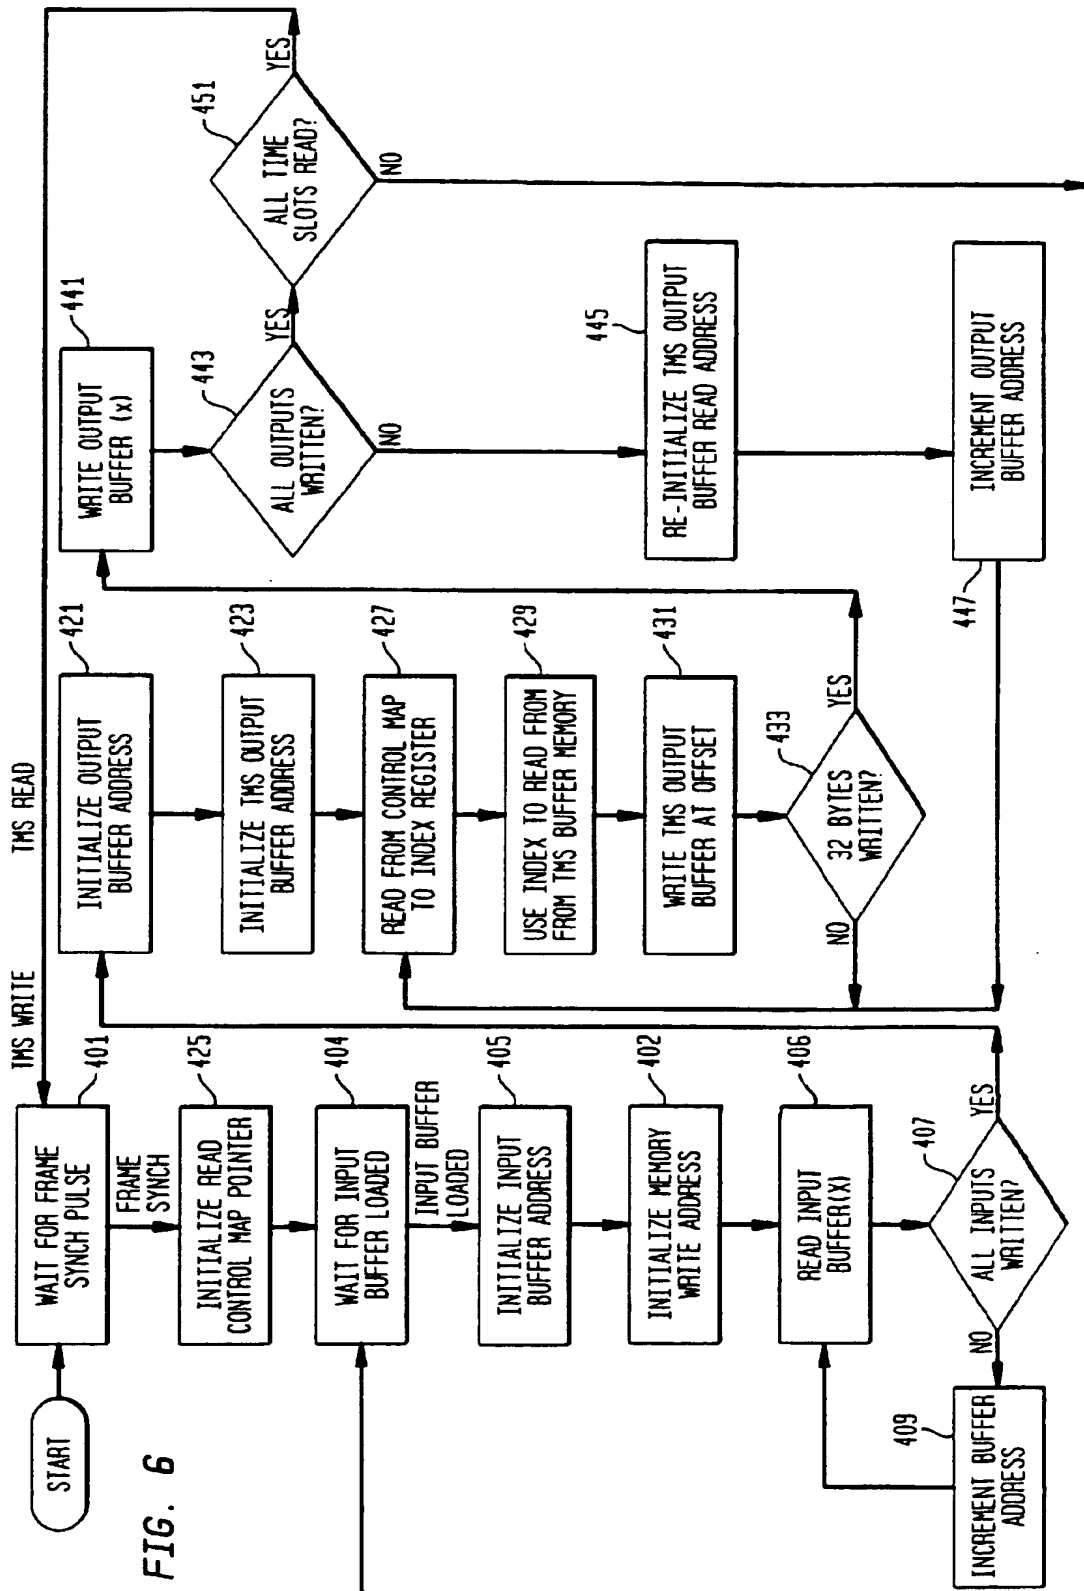


FIG. 3









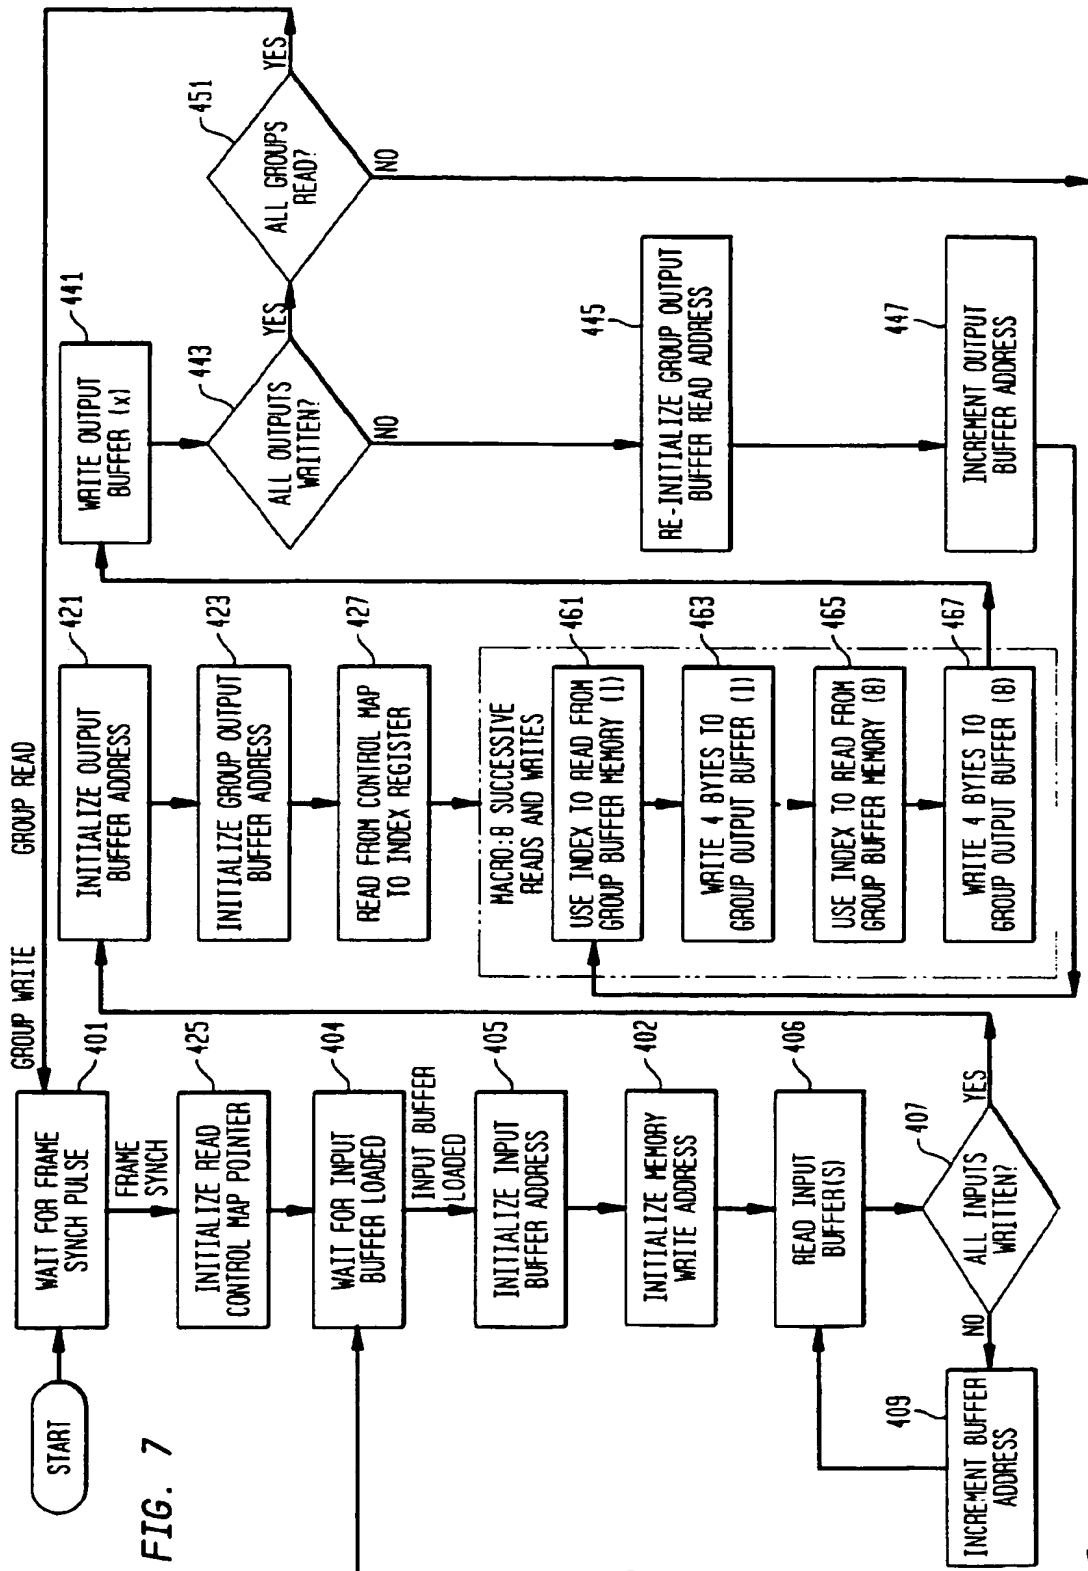


FIG. 8

ATM CELL HEADER
USER NETWORK INTERFACE (UNI)

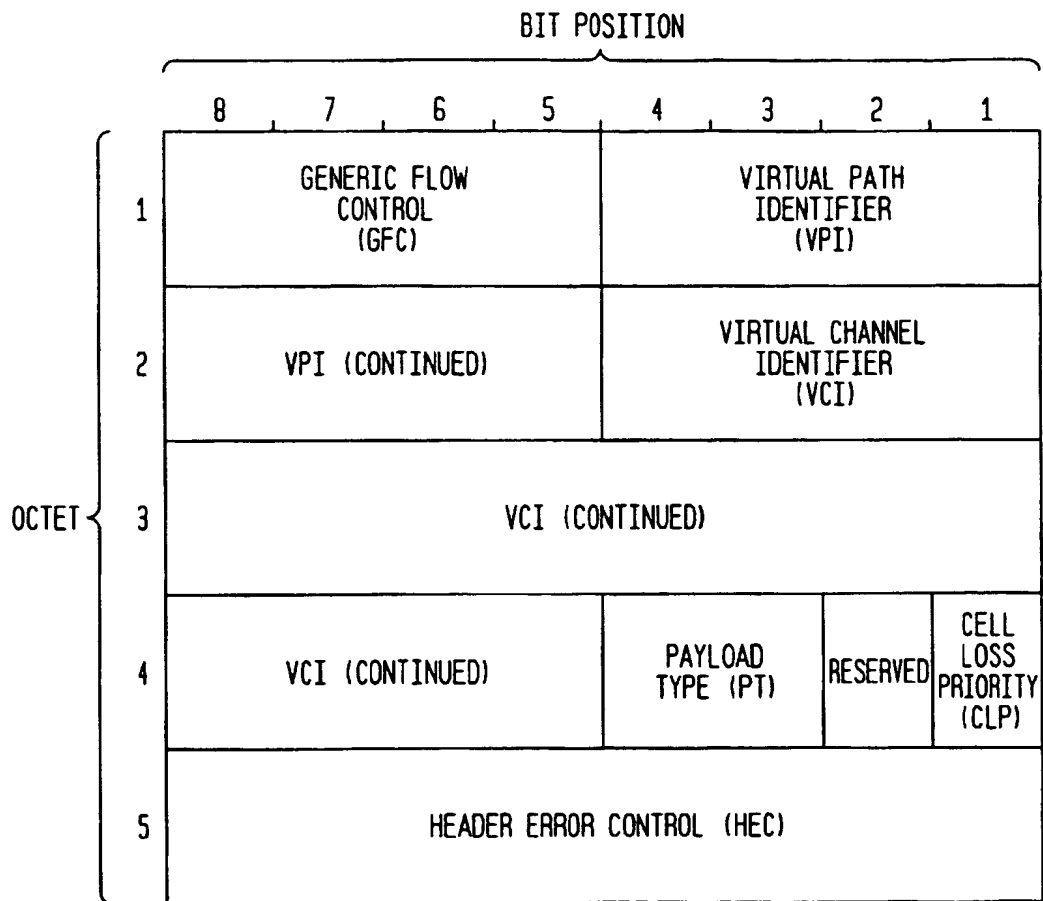


FIG. 9




FIG. 10A




FIG. 10B

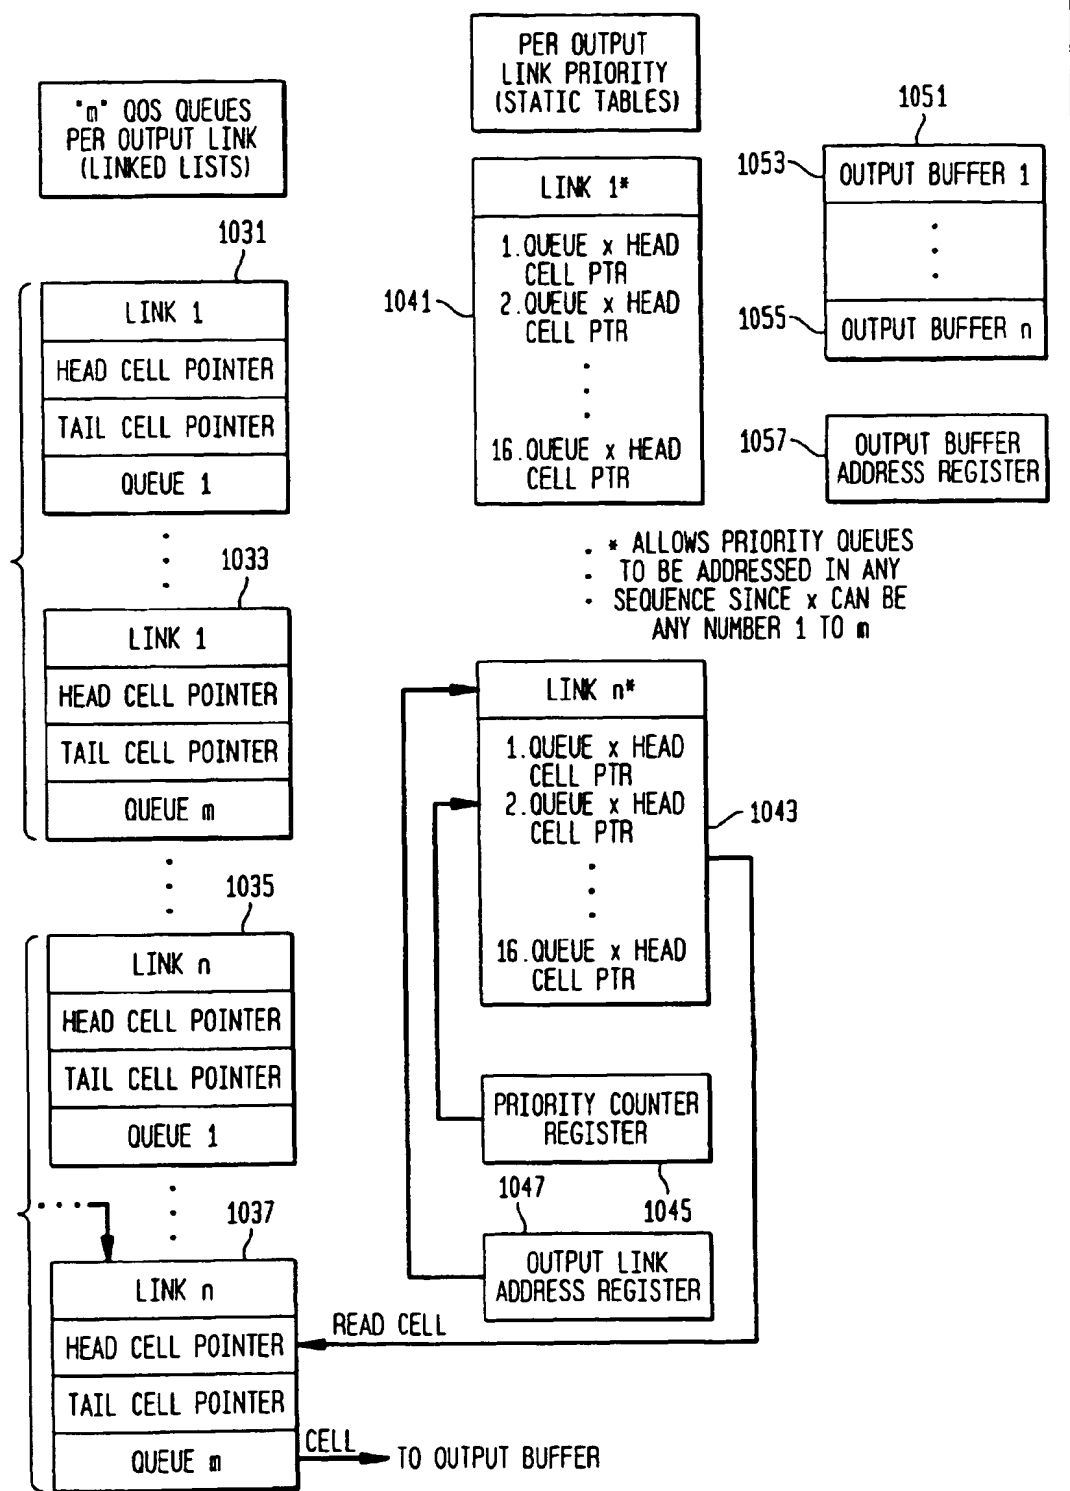


FIG. 11A

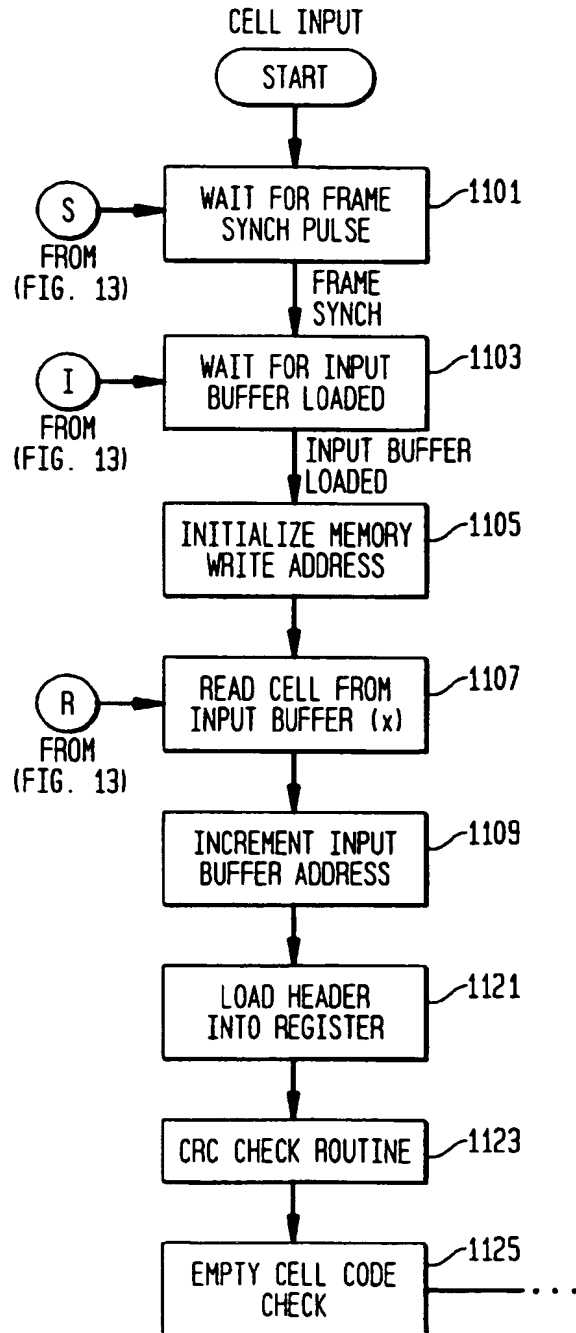


FIG. 11B

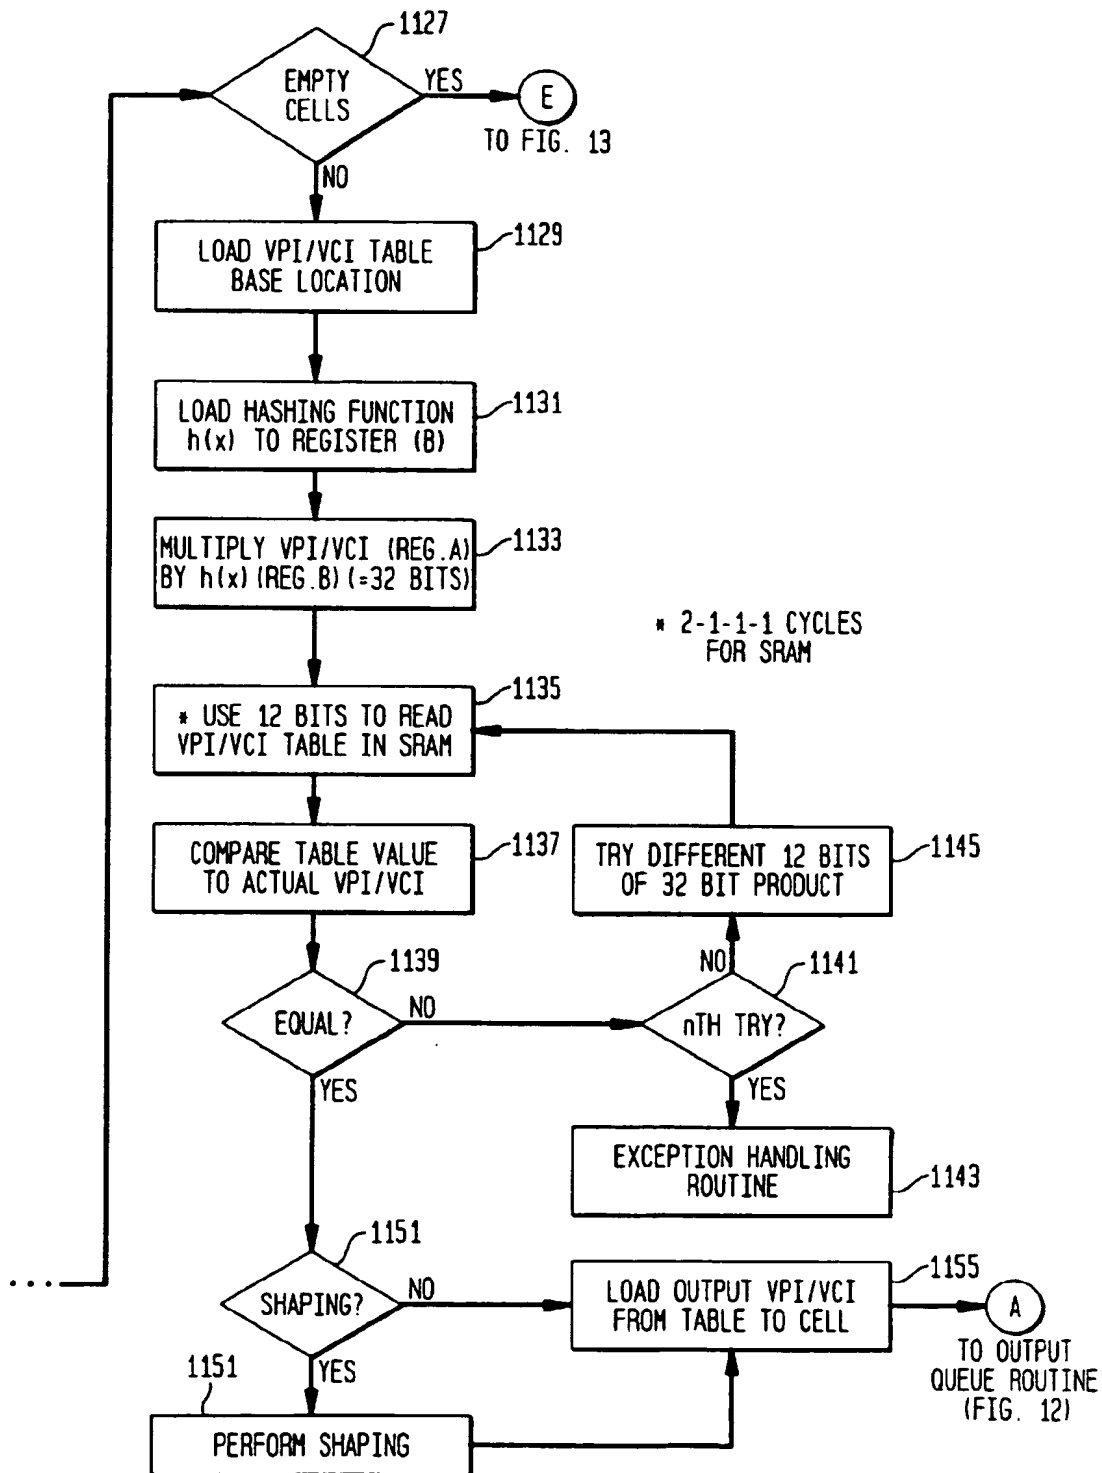


FIG. 12

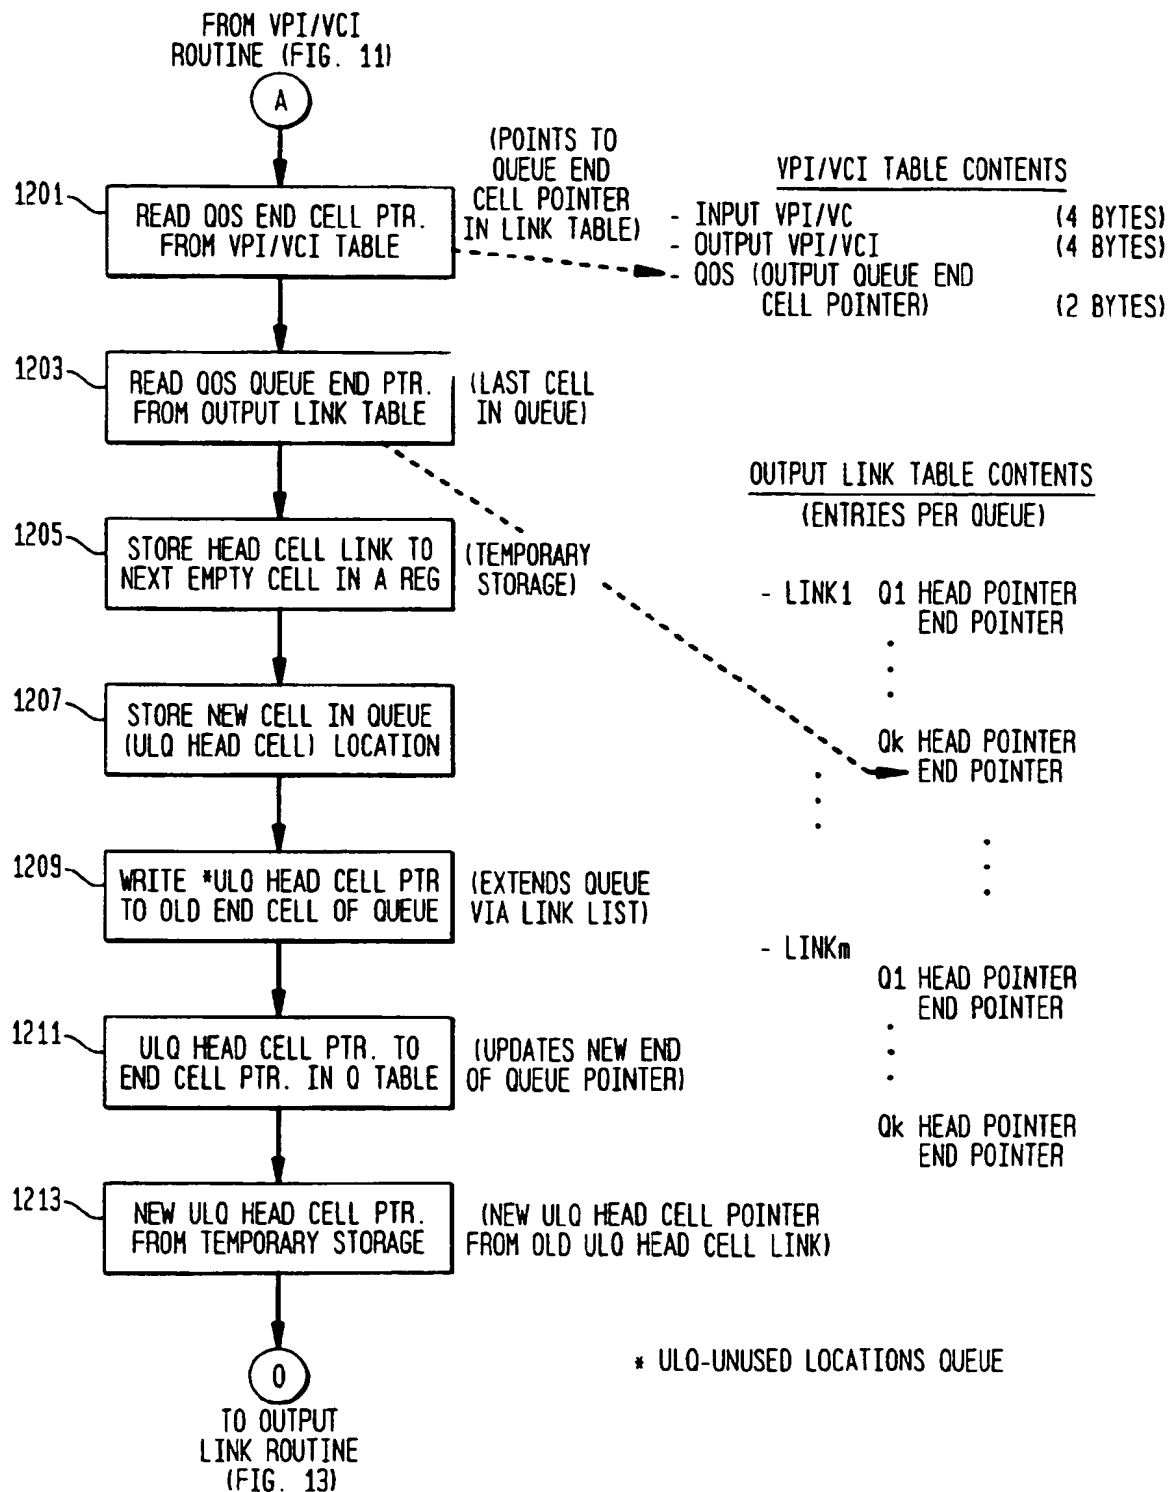


FIG. 13A

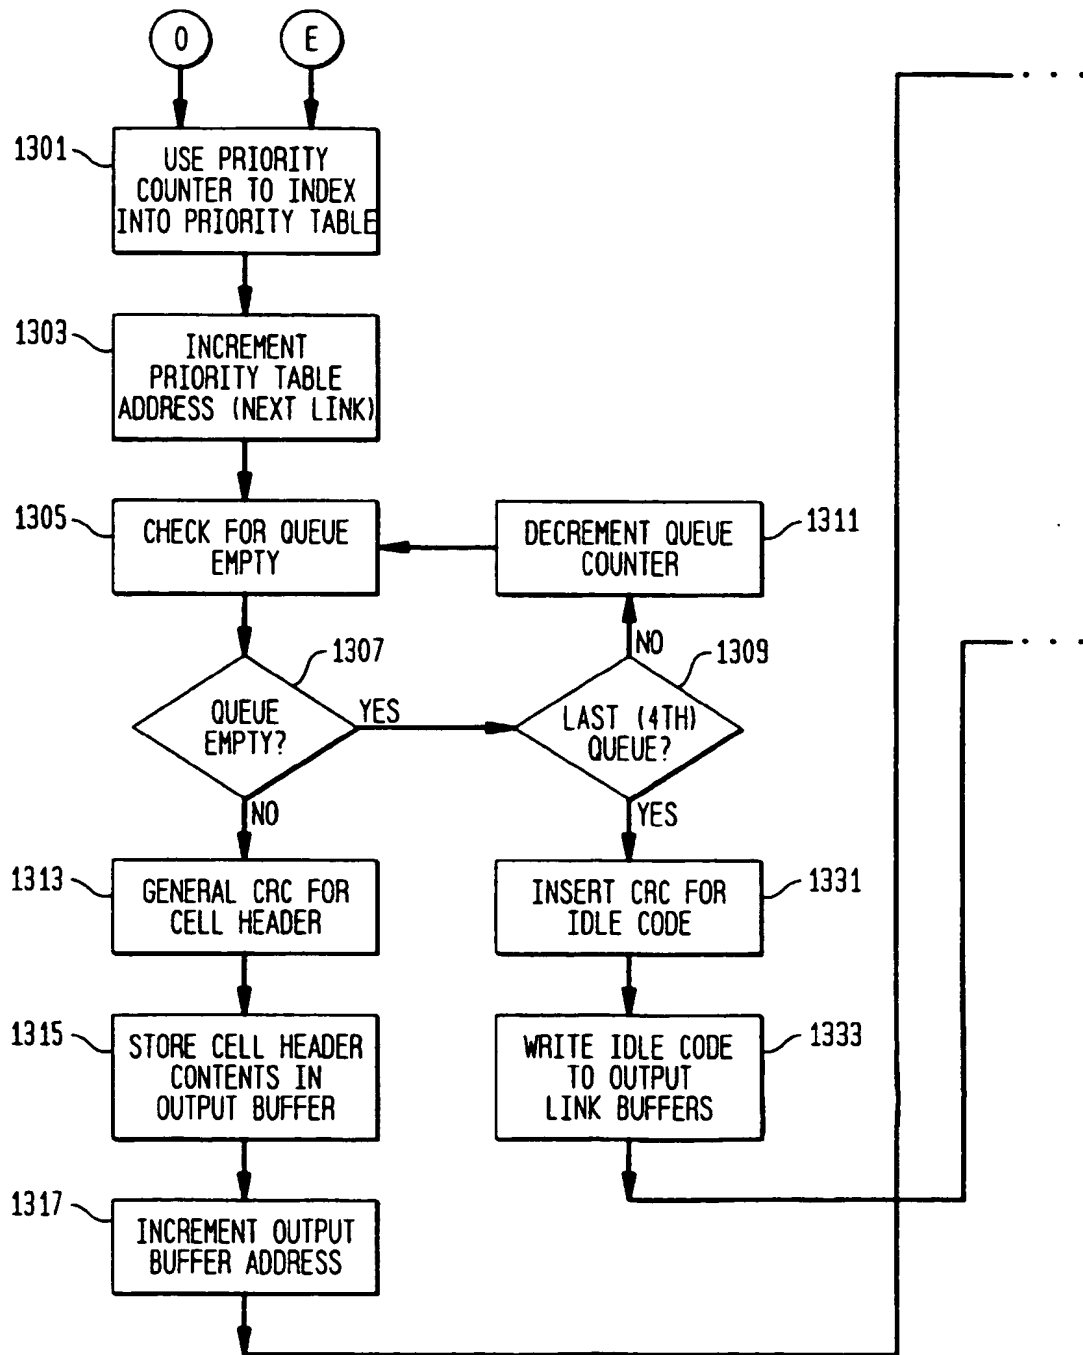


FIG. 13B

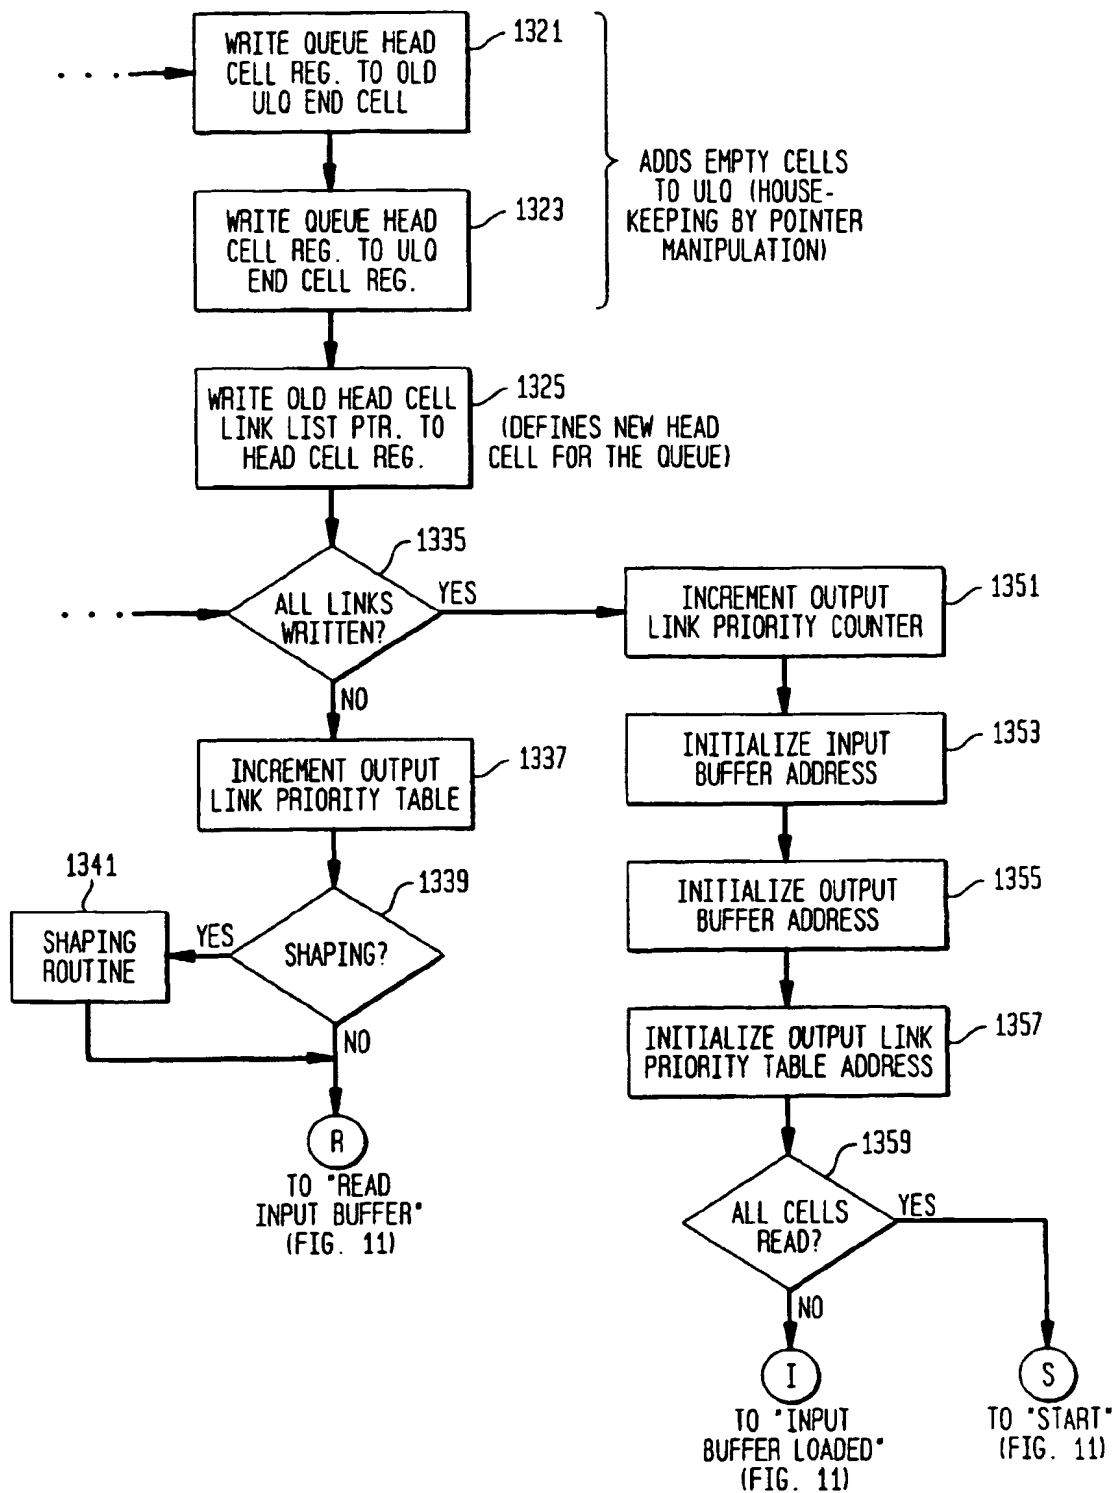


FIG. 14




FIG. 15






FIG. 16





Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 0 987 861 A3**

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:
29.03.2000 Bulletin 2000/13

(51) Int Cl.7: **H04L 12/56, H04Q 11/04,
H04L 12/64**

(43) Date of publication A2:
22.03.2000 Bulletin 2000/12

(21) Application number: **99307100.0**

(22) Date of filing: **07.09.1999**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• **Chrin, Christopher James**
Lisle, Illinois 60532 (US)
• **Zola, Meyer Joseph**
Lisle, Illinois 60532 (US)

(30) Priority: **18.09.1998 US 101030 P**
25.03.1999 US 276187

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr. et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)


(71) Applicant: **LUCENT TECHNOLOGIES INC.**
Murray Hill, New Jersey 07974-0636 (US)

(54) **Flexible telecommunications switching network**

(57) Apparatus, and a method for flexibly switching telecommunication signals. A plurality of input bit streams are connected to a telecommunications switching network fabric element comprising a microprocessor system including memory. The microprocessor system, under program control, performs switching and protocol conversion functions on the input streams in order to generate output streams. Advantageously, a single ele-

ment is able to concurrently switch input signals in a variety of protocols, including circuit switching protocols, such as Pulse code Modulation (PCM), and packet switching protocols, such as Asynchronous Transfer Mode (ATM) protocols and Internet Protocol, (IP). Where desirable, the microprocessor system can also control the protocol conversion of input signals in one protocol to output signals in another protocol.

FIG. 1



EP 0 987 861 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 30 7100

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	EP 0 676 878 A (IBM) 11 October 1995 (1995-10-11)	29,30	H04L12/56 H04Q11/04 H04L12/64
A	* abstract * * page 10, line 31-42 *	1,2	
X	US 5 093 827 A (FRANKLIN ANDREW D ET AL) 3 March 1992 (1992-03-03)	29,30	
A	* abstract * * column 8, line 9-16 *	1,2	
A	US 4 926 416 A (WEIK HARTMUT) 15 May 1990 (1990-05-15)	1,2	
	* abstract * * figures 4,5 *		
A	WO 98 02998 A (HUANG ALAN) 22 January 1998 (1998-01-22)	1,2	
	* abstract * * page 9, line 2-9 *		
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 February 2000	Examiner Dhondt, E
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 (03.82) (F04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 30 7100

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

07-02-2000

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 0676878	A	11-10-1995	CA 2144402 A	15-10-1995
			JP 7288546 A	31-10-1995
			US 5602841 A	11-02-1997
US 5093827	A	03-03-1992	NONE	
US 4926416	A	15-05-1990	DE 3742939 A	06-07-1989
			AT 103129 T	15-04-1994
			AU 2668588 A	22-06-1989
			CA 1330119 A	07-06-1994
			CN 1035927 A, B	27-09-1989
			DE 3888480 D	21-04-1994
			EP 0320772 A	21-06-1989
			ES 2052679 T	16-07-1994
			JP 1998813 C	08-12-1995
			JP 2004054 A	09-01-1990
			JP 7028314 B	29-03-1995
			KR 9607670 B	08-06-1996
			MX 174401 B	13-05-1994
			NO 173680 C	12-01-1994
WO 9802998	A	22-01-1998	US 5841775 A	24-11-1998

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office No. 12/82